

How Homophily Affects Communication in Networks*

Benjamin Golub[†]

Matthew O. Jackson[‡]

December 23, 2008

Abstract

We examine how three different communication processes operating through social networks are affected by homophily – the tendency of individuals to associate with others similar to themselves. Homophily has no effect if messages are broadcast or sent via shortest paths; only network connection density matters. In contrast, homophily substantially slows learning based on repeated averaging of neighbors’ information and Markovian random walks such as the Google random surfer model. Indeed, the latter processes are strongly affected by homophily but completely independent of connection density, provided this density exceeds a low threshold. We obtain these results by establishing new results on the spectra of large random graphs and relating the spectra to homophily. We conclude by checking the theoretical predictions using observed high school friendship networks from the Adolescent Health dataset.

Keywords: Networks, Homophily, Random Graphs, Friendships, Social Networks, Learning, Mixing Time, Convergence, Speed of Learning, Speed of Convergence

JEL Classification Numbers: A14, D85, I21, J15

*Jackson gratefully acknowledges financial support from the NSF under grant SES-0647867. 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 James Moody for making available the network data from the ‘Add Health’ data set. We also thank Christoph Kuzmics for comments on an early presentation.

[†]Graduate School of Business, Stanford University. Email: bgolub@stanford.edu, <http://www.stanford.edu/~bgolub/>.

[‡]Department of Economics, Stanford University and external faculty of the Santa Fe Institute. Email: jacksonm@stanford.edu, <http://www.stanford.edu/~jacksonm/>

1 Introduction

How does a society’s structure affect the speed at which information diffuses within it? In particular, how do the segregation patterns in a social network affect how information is diffused and aggregated within that society? How does that relationship change as we vary the communication process? In this paper, we take a step toward answering these questions by studying how information transmission is affected by one of the most salient observed features of social networks – homophily.

Homophily – the tendency of individuals to associate with those similar to themselves – has been observed since antiquity¹ and studied by sociologists under that name since Lazarsfeld and Merton (1954). It has been documented across a wide array of different characteristics, including race, age, ethnicity, profession, religion, and various behaviors (see McPherson, Smith-Lovin and Cook (2001) for a survey). Indeed, homophily is one of the most pervasive and robust tendencies of social networks, and one of the main features that makes them distinctively social. Nevertheless, there has been effectively no modeling of the impact of homophily on important processes such as the diffusion of information, the formation of consensus, or other forms of social learning and communication that are integral to social and economic behaviors.

In this paper, we address this gap by modeling both homophily and communication explicitly at the level of interactions between individual agents, and then using the model to examine how the extent of homophily affects communication in various settings. In doing this, we bring together two important but previously separate threads of research: statistical models of networks with homophily and models of information diffusion and learning in a given network. Random network modeling allows us to capture homophily in a natural way: certain links are more likely to form than others but the precise realization of the network is random. Working with models of communication and learning in social networks allows us to understand how these are affected by changes in homophily.

We use one probabilistic model of homophily: the *multi-type random network*. This is a generalization of a seminal Erdős-Rényi random network model, and it also nests many other previous random network models. In this model, agents are divided into different types, and the probability that a link is formed depends only on the types of the agents involved. Using these random networks as a basis, we study three different learning processes². Whether or

¹By Plato’s time, homophily was already considered proverbial. In *The Republic*, Cephalus says: “For it often happens that some of us elders of about the same age come together and verify the old saw of like to like” (Book I, p. 329).

²There are some similar models that generalize the simpler random network models, such as variations on

not homophily affects the speed of communication or learning at all turns out to depend on how information is diffused through the network. The first process we study is one in which information is either broadcast or navigated to its destination via shortest paths. This class includes many peer-to-peer systems like the Internet, mechanisms where messages are routed within an organization using something like an organizational chart, and information spreading phenomena where people tell an important piece of news to everyone they know. The second is a process based on a linear updating or learning model developed by French (1956) and Harary (1959) where individuals update their beliefs or actions by repeatedly taking weighted averages of their neighbors' beliefs or actions. This captures boundedly rational processes of updating as well as pressures to conform or a desire to match actions of neighbors³. The third is a random walk process on a network, where some particle hops around the network, having equal probability of moving along any link out of its current node; one example is Google's famous model of a surfer who randomly follows a link out of the website he is currently visiting. These three processes encompass many important forms of network-based communication and diffusion.

Our main results show that whether or not homophily has an impact depends on the communication process, and then we also detail precisely how homophily matters when it matters. In particular, we show that processes based on shortest paths are unaffected by homophily, while averaging processes and random walks are affected and can be substantially slowed down by homophily. A central intuition behind these results is that homophily does not affect the average distance between nodes in a network if connectivity is held constant, but it does affect the relative numbers of paths between various pairs of agents. The reason that it does not affect average shortest path length is that even with substantial homophily and the resulting clustering, extended neighborhoods still expand exponentially in the number of steps. Thus, even though there may be fewer paths between agents of different types, the average shortest path length between them does not change substantially. So processes in which information must simply traverse a shortest path have essentially the same behavior as homophily changes. In contrast, processes based on weighted averaging or random walks are substantially slowed down as homophily increases. This is because even though the average path length is unchanged, there are relatively fewer paths between agents of different types as homophily increases. This means that nodes are more influenced

p^* models (e.g., see the discussion in Jackson (2008b)). However, those have generally been used as empirical models to fit observed patterns of relationships, and have not been used to examine processes occurring on networks, as we do here.

³Such updating processes can lead a society to an optimal aggregation of information in some settings, depending on the specifics of the social network structure (e.g., see Golub and Jackson (2007)).

by others of their own type, which reinforces global heterogeneity in beliefs or behaviors and slows down convergence to a steady state.

An interesting implication of our results is the following. When looking at shortest path-based processes, increasing homophily does not affect the average shortest path length. In contrast, with averaging and random-walk based processes, if homophily is increased, the speed of convergence decreases even if link density also goes up. That is, these processes are completely orthogonal to each other when it comes to what matters for their speeds.

As an empirical illustration of the results, we examine the time to convergence of these processes on social networks from over eighty different high school friendship networks that exhibit varying degrees of homophily. We show that the predictions of our analysis fit well and that the speed of convergence depends on homophily in the predicted ways.

Conceptual Outline

Our results involve several layers, including some contributions on the mathematical side which are needed to deduce the relationships discussed above. In particular, in order to relate network structure (including homophily) to convergence speeds we have to work with the spectral decomposition of a network, which naturally leads us to develop new results on the spectra of large random graphs. Given this layering, it is useful to have a road map of how all of our results fit together.

After introducing the model and background definitions, we first present the main conceptual results that relate homophily to speed of communication. Having the main conclusions in hand, we then present a series of results that are used to derive those conclusions. In particular, we begin with a statement of standard results relating speed of convergence to second eigenvalues. Next, we present our key technical theorem, showing that the second eigenvalue associated with a random network will be close to the second eigenvalue of a smaller matrix which deals only with relative linking probabilities across types. That is, all that really matters in determining the second eigenvalue in large societies is the expected connection probabilities between types. This result allows us to derive second eigenvalues for the multi-type random networks based on homophily patterns and relate the eigenvalues to simple measures of homophily. Thus, the key technical result allows us to tie homophily to second eigenvalues, which in turn govern consensus and mixing times. Putting all of this together provides our conceptual conclusions.

2 The Model: Networks and Processes

2.1 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 we restrict attention to undirected networks.⁴

Let $d_i(\mathbf{A}) = \sum_{j=1}^n A_{ij}$ denote the degree of node i . Let $d_{\min}(\mathbf{A})$ and $d_{\max}(\mathbf{A})$ be the minimum and maximum degrees, respectively, and $\bar{d}(\mathbf{A})$ denote average degree, and let

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

be the total degree in the society.

2.2 Multi-Type Random Networks

In order to study the impact of homophily on communication and learning through a network, we introduce a random network model that incorporates homophily. The seminal random network model of Erdős-Rényi random networks is a special case of the model here (and the same is true of the model based on degree distributions of Chung and Lu (2002)) which allows us to make benchmark comparisons to the literature on speeds of processes on networks without homophily.

The structure we use to model homophily is what we call the *multi-type random network*. It consists of a vector $\mathbf{n} = (n_1, \dots, n_m)$ which captures how many nodes of each type there are (and implicitly, how many types, m , there are), and a symmetric m -by- m matrix \mathbf{P} , whose entries in $[0, 1]$ describe the probabilities of links between various types. Let N_k be the set of nodes of type k , and without loss of generality label nodes so that $\{1, \dots, n_1\}$ are the nodes of the first type, $\{1 + n_1, \dots, n_1 + n_2\}$ are the nodes of the second type, and $N_k = \{1 + \sum_{i < k} n_i, \dots, \sum_{i \leq k} n_i\}$ are the nodes of the k -th type. The resulting random network is captured via its adjacency matrix which is denoted $\mathbf{A}(\mathbf{P}, \mathbf{n})$ and is a random variable. In particular, $\mathbf{A}(\mathbf{P}, \mathbf{n})$ is built by letting the entries A_{ij} with $i > j$ be independent Bernoulli random variables with parameter $P_{k\ell}$ if $i \in N_k$ and $j \in N_\ell$. That is, the entry $P_{k\ell}$

⁴Although we conjecture that the results can be extended to directed networks without much change in the statements (as the communication/learning processes have direct extensions to the directed case), there are parts of the proofs that take advantage of the symmetry of the adjacency matrix, and so we are not sure of what modifications would ensue in examining directed networks.

captures the probability that an agent of type k links to an agent of type ℓ . We then fill in the remaining entries of \mathbf{A} by symmetry: $A_{ij} = A_{ji}$.

Here are some special cases of the model.

If $P_{k\ell} = p$ for all k, ℓ , then this is simply an Erdős-Rényi random network.

The random network model of Chung and Lu (2002) is the special case where the only heterogeneity in liking is induced by expected degrees. In particular, each type has an expected degree w_k , and $P_{k\ell} = w_k w_\ell / W$ where $W = \sum_k n_k w_k$. Thus, the matrix \mathbb{P} is reduced from having m^2 degrees of freedom to having m , namely the expected degrees of the nodes.

A spatial model is one where each type of node has a parameter $\theta_k \in \mathbb{R}^m$ for some m that describes it, and $P_{k\ell} = f(\Delta(\theta_k, \theta_\ell))$ where f is a decreasing function and Δ is Euclidean distance. This the probability that nodes link to each other is a function of how similar their types are.

2.2.1 The Islands Model

Another special case of the model that we discuss in some of the results below is an *islands model*. There we assume that all n_i are equal, so that islands are equal-sized; we set $P_{kk} = p_s$ for all k and $P_{k\ell} = p_d$ for all $k \neq \ell$. Thus, nodes of the same type connect to each other with one probability, and nodes of different types connect to each other with another probability.⁵

These examples are only a few of the possibilities, and clearly one can consider combinations of these variations, and other considerations such as special cases where linking probabilities are built on some hierarchy, etc.

2.2.2 Remarks on When the Multi-Type Random Networks Model is Useful

These examples give an idea of how rich the multi-type random networks model is. However, as with any model, it is most pointed in its predictions when we obtain a significant reduction in the dimensionality of the problem. In particular, for our main results involving representative agents to be most useful, it is helpful for there not to be too many types or, failing that, for the interaction between types to be described by only a few parameters. Effectively, the results reduce the problem of working with a network of n individuals to a simpler problem of working with m types. If there are as many types as individuals, then

⁵See Currarini, Jackson and Pin (2009) and Copic, Jackson and Kirman (2005) for illustrations and applications of such a model.

clearly that will be unhelpful without additional assumptions; however, a good deal of explanatory power comes out from looking at just a few types, to the extent that a few types capture most of the important variation in the data. As we will see from our look at the data, very simple definitions of types have substantial explanatory power.

2.3 Communication and Learning Processes

We now describe three different processes communication/learning that we consider. As we shall see, there will be some differences in how these are affected by homophily.

2.3.1 Shortest-Path Communication

Consider any of a class of processes where the time it takes to communicate between two nodes is proportional to the shortest path between the two nodes.⁶ In a connected network, this applies to broadcast processes, where nodes communicate to all neighboring nodes in each period, or to processes where the network is explicitly navigated by a traveler using some sort of addressing system. This applies to some social and many physical transmission processes.

2.3.2 A Repeated Updating Learning Model

The second process that we examine is based on a model first discussed by French (1956) and Harary (1959), and articulated in a more general form by DeGroot (1974).

Given a network \mathbf{A} , let $\mathbf{T}(\mathbf{A})$ be defined by $T_{ij}(\mathbf{A}) = A_{ij}/d_i$. Beginning with some initial belief vector $\mathbf{b}(0) \in [0, 1]^n$, let

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

for all $t \geq 1$. That is, agents form today's beliefs by taking the average of neighbors' beliefs yesterday, where an agent can be his own neighbor. It is immediate that then

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

If the initial beliefs $\mathbf{b}(0)$ 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 with certain priors about signal precisions as discussed by DeMarzo, Vayanos, and Zwiebel (2003). The behavioral aspect of the model concerns times after

⁶Standard network definitions, such as shortest path, are omitted. See Jackson (2008b) for background definitions.

the first round of updating. Here, it is no longer Bayesian to update using a weighted-average rule, but due to the overwhelming complexity of the Bayesian calculation, we assume agents continue using the simple averaging rule in later periods, too. More discussion of this assumption can be found in DeMarzo, Vayanos, and Zwiebel (2003).

Beyond the interpretation of updating signals and “learning”, the linear updating model can also be interpreted as a model that captures behaviors where agents adjust their behaviors to match the average of their neighbors’ choices. In particular, it can be interpreted as myopic best-response updating in a game. Suppose that agents have to choose each period a variable b_i , which captures their behaviors, for instance which dialect of a language they speak, and the dialects correspond to points in $[0, 1]$. The cost to i of communicating with j is $(b_i - b_j)^2$. If each agent communicates with his neighbors according to \mathbf{A} , then the best response mapping is given by the linear updating rule.

If $\mathbf{T}(\mathbf{A})$ is not connected, then it suffices to consider the asymptotic behavior of each connected component to understand the full dynamics of the process.⁷ Thus, we assume from now on that $\mathbf{T}(\mathbf{A})$ is connected, and define $\mathbf{T}(\mathbf{A})^\infty = \lim_{t \rightarrow \infty} \mathbf{T}(\mathbf{A})^t$.

LEMMA 1. *If \mathbf{A} is connected, then $\mathbf{T}(\mathbf{A})^t$ converges to a limit $\mathbf{T}(\mathbf{A})^\infty$ such that $(\mathbf{T}(\mathbf{A})^\infty)_{ij} = \frac{d_i(\mathbf{A})}{D(\mathbf{A})}$.*

Lemma 1 follows from standard results on Markov chains (e.g., see Golub and Jackson (2007) and Jackson (2008b) for details and background) and implies that for any given initial vector of beliefs $\mathbf{b}(0)$, the limiting belief

$$\lim_t \mathbf{b}(t) = \mathbf{T}^\infty \mathbf{b}(0) = (b, b, \dots, b) \quad \text{where} \quad b = \sum_i \frac{b_i(0) d_i(\mathbf{A})}{D(\mathbf{A})}.$$

Thus, the relative influence that an agent has over the final beliefs is his or her relative degree.

2.3.3 Random Walks

The third process that we study is a random walk on a network. This is a process where a particle starts at some node and hops to any of its neighbors with equal probability at each step. One example to think of is that of a college student who is viewing Facebook profiles; at each step, she clicks on a random friend of the person whose profile she is currently

⁷If the communication network is directed then convergence requires some aperiodicity in the cycles of the network and works with a different segmentation into components, but still holds quite generally, as discussed in Golub and Jackson (2007).

viewing. Another example is the Google model of a surfer who randomly clicks on links as he navigates the World Wide Web.⁸ Here, the particle starts at some location and transitions from node i to node j with probability T_{ij} . The question is how long it takes to reach the steady state distribution on location. While this is a fairly specific process, and may not capture as many applications as the previous two processes, it has figured prominently in the literature Markov processes and random graphs and so it is a very useful benchmark.

Just as in the case of the linear updating learning model, Lemma 1 implies that if \mathbf{A} is connected, then \mathbf{T}^t converges and then the limit distribution of the random walk is to be at node i with probability $\frac{d_i(\mathbf{A})}{D(\mathbf{A})}$ regardless of the starting position of the random walk. So, the limiting distribution of the time that the walk spends at a given node is proportional to its degree.

2.4 Consensus Time and Mixing Time

We now present ways of measuring the speed at which the above-defined processes operate on a given network. The different processes suggest different measures.

First, shortest-path based processes have an obvious measure of speed, which is simply the average shortest path length in the network, or 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. These are standard notions, so there is no need to develop any special measure for such processes.

The other two processes require measures of timing/distance that are more tailored to them. We now discuss each in turn.

2.4.1 Distances between Vectors

There are two distance measures that we focus on in measuring convergence.

The first is a standard weighted squared deviation distance. Given two vectors of beliefs v and u let

$$\|\mathbf{v} - \mathbf{u}\|_{\mathbf{w}}^2 = \sum_i w_i (v_i - u_i)^2.$$

In applying this, we will be interested in the differences between beliefs at time t and

⁸Of course, we might think of him being biased toward following certain out-links from a given web page; this could be modeled by using a nonuniform random walk; i.e., one in which all nonzero entries in a given row of \mathbf{T} are not the same. We suspect that our conclusions would simply be modified by weighting factors, but the symmetry of the simpler case is handy in our proofs.

their limit:

$$\|\mathbf{T}(\mathbf{A})^t \mathbf{b} - \mathbf{T}(\mathbf{A})^\infty \mathbf{b}\|_{\mathbf{w}}^2.$$

It will be useful to use weights, \mathbf{w} , that are the influences of the agents $\mathbf{s}(\mathbf{A})$, where

$$\mathbf{s}(\mathbf{A}) = \left(\frac{d_1(\mathbf{A})}{D(\mathbf{A})}, \dots, \frac{d_n(\mathbf{A})}{D(\mathbf{A})} \right).$$

The distance

$$\|\mathbf{T}(\mathbf{A})^t \mathbf{b} - \mathbf{T}(\mathbf{A})^\infty \mathbf{b}\|_{\mathbf{s}(\mathbf{A})}^2$$

examines the squared difference between agents' current beliefs and their limit beliefs. The distance is weighted by the agents' degrees which gives more weight to relatively more influential agents. This quantity has a fairly simple interpretation. Consider the following experiment: agents start with beliefs \mathbf{b} ; at time t , an agent is sampled uniformly at random. We imagine that he asks a random neighbor for his opinion: i.e. one of his neighbors is sampled uniformly at random; we record the square of this neighbor's deviation from consensus beliefs. The expectation of this variable under this experiment is the distance defined above. In other words, sampling each agent in proportion to his degree captures the deviation from consensus of an opinion sent at time t across a randomly chosen link in the network.

The following lemma shows an obvious relationship between a straight averaging of the squared deviations in beliefs and this weighted averaging.

LEMMA 2. *Let $\mathbf{e} = (1, \dots, 1)$. Then*

$$\begin{aligned} \frac{\bar{d}(\mathbf{A})}{d_{\max}(\mathbf{A})} \|\mathbf{T}(\mathbf{A})^t \mathbf{b} - \mathbf{T}(\mathbf{A})^\infty \mathbf{b}\|_{\mathbf{s}(\mathbf{A})}^2 &\leq \|\mathbf{T}(\mathbf{A})^t \mathbf{b} - \mathbf{T}(\mathbf{A})^\infty \mathbf{b}\|_{\mathbf{e}/n}^2 \\ &\leq \frac{\bar{d}(\mathbf{A})}{d_{\min}(\mathbf{A})} \|\mathbf{T}(\mathbf{A})^t \mathbf{b} - \mathbf{T}(\mathbf{A})^\infty \mathbf{b}\|_{\mathbf{s}(\mathbf{A})}^2. \end{aligned}$$

Thus, for graphs where the highest- and lowest-degree agents have degrees not too different from the average, as in some of the networks we will be concerned with, whether or not we weight mean square deviation from consensus by degree will not make a large difference. Even in networks where there are large deviations in degree, if the number of deviant nodes is bounded, then a direct variation of Lemma 2 implies that the two notions are still close. Given that converting between these two notions only requires computing some bounds which will usually be good but which will depend on the application, we work with the degree-weighted version of the deviation measure as it has nicer mathematical properties and more intuitive relationships to the network structure.

The other distance measure that we will work with is the total variation metric, where the distance between a vector \mathbf{v} and another vector \mathbf{u} is

$$\|\mathbf{v} - \mathbf{u}\|^{TV} = \frac{1}{2} \sum_i |v_i - u_i|.$$

We will be applying this in cases where \mathbf{v} and \mathbf{u} are probability measures, so that $\mathbf{v} \geq \mathbf{0}$, $\mathbf{u} \geq \mathbf{0}$, and $\sum_i v_i = \sum_i u_i = 1$. Then it is straightforward to see that

$$\|\mathbf{v} - \mathbf{u}\|^{TV} = \frac{1}{2} \sum_i |v_i - u_i| = \max_{C \subset N} \left| \sum_{i \in C} v_i - \sum_{i \in C} u_i \right|,$$

so the total variation metric keeps track of the maximal difference in the probability that two measures assign to some set.

2.4.2 Consensus Time and the Linear Updating Model

A central question in the linear updating model is the rate at which beliefs of the society converge to their consensus limit. We define the consensus time as follows.

DEFINITION 1. *The consensus time to $\varepsilon > 0$ of the network \mathbf{A} is*

$$\text{CT}(\varepsilon; \mathbf{A}) = \sup_{\mathbf{b} \in [0,1]^n} \min\{t : \|\mathbf{T}(\mathbf{A})^t \mathbf{b} - \mathbf{T}(\mathbf{A})^\infty \mathbf{b}\|_{\mathbf{s}(\mathbf{A})}^2 < \varepsilon\}.$$

The need to consider different potential starting belief vectors \mathbf{b} is clear, as if one starts with $b_i(0) = b_j(0)$ for all i and j then consensus is reached instantly. Thus, the “worst case” \mathbf{b} will generally have 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 (e.g., see Montenegro and Tetali (2006)).

Since \mathbf{T} is a contraction under the distance measure (a standard fact about reversible Markov chains), once the mean-square deviation is below ε , it can never go above it again. Thus, the definition is equivalent to letting $\text{CT}(\varepsilon; \mathbf{A})$ be the earliest time such that deviation from consensus is small forever after.

2.4.3 Mixing Time and Random Walks

There are many definitions measuring the distance of a random process to its limit that come from the literature on Markov chains (e.g., see Montenegro and Tetali (2006)), and the

following definition is among the most common. Let \mathbf{e}_i be the unit vector with an entry of 1 in the i -th entry and 0's elsewhere.

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

$$\text{MT}(\varepsilon; \mathbf{A}) = \sup_i \min\{t : \|\mathbf{e}_i \mathbf{T}(\mathbf{A})^t - \mathbf{e}_i \mathbf{T}(\mathbf{A})^\infty\|^{TV} < \varepsilon\}.$$

Mixing time keeps track of how different the probability across states is after t periods compared to the limiting distribution. Taking the supremum over different starting states is equivalent to considering all possible starting distributions.

2.4.4 The Relation of Consensus and Mixing Time

The consensus and mixing times have a close relationship intuitively, as both depend on how quickly \mathbf{T}^t approaches its limit. The main difference is that consensus time works with mean-squared deviations while mixing time works with the sum of absolute values of the deviations. In terms of the mathematics, the difference between them is simply the difference between the ℓ^2 and ℓ^1 norms as well as a difference in the normalizing constant.

Just as an illustration, consider the distance between $(1, 0, 0, \dots)$ and $(1/n, 1/n, 1/n, \dots)$. If we think of these as probability measures, then they are quite different as one is a Dirac measure and the other is a uniform distribution. In contrast, if these are behaviors or beliefs, then only one agent in the society is deviating substantially from the limiting behavior or beliefs. Thus depending on the application, one might or might not want to consider these to be close or far apart. Under the ℓ^1 norm used in calculating consensus time, these are close to each other, while under the ℓ^2 norm used in calculating mixing time, these are quite far apart.

The mixing time approach is most natural and standard in the setting of Markov chains, where distributions are important, and consensus time is a natural measure in the setting of linear updating models. We present results on both, and shall see that despite the differences, they will behave similarly in our setting under the right normalizations. As we shall see also in the empirical section (see Figure 6), consensus time and mixing time will essentially coincide in the high school data.

2.4.5 Asymptotics

In some cases, we consider what happens as n grows. This is natural as there are many properties of random graphs that can be deduced to hold almost surely for “large” networks,

but that are hard to express in any meaningful way for a small random graph where any possible configuration has a nontrivial probability of arising. In such cases, then there is some question as to the appropriate choice of ε .

Given the steady state distribution $\mathbf{s}(\mathbf{A}) = \left(\frac{d_1(\mathbf{A})}{D(\mathbf{A})}, \dots, \frac{d_n(\mathbf{A})}{D(\mathbf{A})} \right)$, if the $d_i(\mathbf{A})$'s are growing at roughly the same speed (a condition in some of the results below), then the entries of $\mathbf{s}(\mathbf{A})$ will be of order $1/n$. As such, a natural benchmark is to examine $\text{CT}(\gamma/n^2; \mathbf{A})$ (given the squaring in the norm) and $\text{MT}(\gamma/n; \mathbf{A})$ for some fixed $\gamma > 0$.

3 The Speed of Communication

We now present our main conceptual conclusions about the speed of communication and learning for the various processes and discuss the contrasts across different sorts of communication. We then come back to the main technical contributions in the next section, which seem to be of some independent interest.

3.1 Shortest-Path Communication

Consider the multi-type random network (\mathbf{n}, \mathbf{P}) with an associated number of nodes n and let $d_k(\mathbf{n}, \mathbf{P}) = \sum_{k'} P_{kk'} n_{k'}$ indicate the expected degree of a node in group k , $\tilde{d}(\mathbf{n}, \mathbf{P}) = \sum_k (d_k(\mathbf{n}, \mathbf{P}))^2 n_k / D(\mathbf{n}, \mathbf{P})$ be the second order average degree in society,⁹

$$D(\mathbf{n}, \mathbf{P}) = \sum_k d_k(\mathbf{n}, \mathbf{P}) n_k$$

be the total expected degree, and $p(\mathbf{n}, \mathbf{P}) = D(\mathbf{n}, \mathbf{P})/n$ be the average probability of a link. Suppose that

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

These conditions admit many cases of interest and can be understood as follows: (i) implies that there is not a divergence in the expected degree across groups; (ii) ensures that

⁹Note that if the average degree $d_k(\mathbf{n}, \mathbf{P})$ is the same across groups, then this is just the average degree.

the average degree grows with n fast enough so that the network becomes connected with a probability going to 1, so that the network will not have isolated components across which communication is impossible; (iii) implies that average degree grows more slowly than n , as otherwise the shortest path degenerates to being of length 1 or 2 and does not match most empirical applications; and (iv) 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. These conditions still allow for substantial homophily. For instance, if n_k is on the order of $d(\mathbf{n}, \mathbf{P})$, then this still allows the probability of a link within same group to even become infinite relative to the probability across links, so that arbitrarily high levels of homophily are permitted.

THEOREM 1 (Jackson (2008a)). ¹⁰ *If the random network process (\mathbf{n}, \mathbf{P}) satisfies (i)-(iv) then, asymptotically almost surely in n , $\mathbf{A}(\mathbf{n}, \mathbf{P})$ is connected; the average distance between nodes is $(1 + o(1)) \log(n) / \log(\tilde{d}(\mathbf{n}, \mathbf{P}))$; and the diameter of the largest component is $\Theta(\log(n) / \log(\tilde{d}(\mathbf{n}, \mathbf{P})))$.*

Theorem 1 tells us that although homophily can change the basic structure of a network, it does not affect the average shortest-path distance between nodes in the network. Moreover, we have a precise expression for that average distance which is the same as it is in an Erdős-Rényi random network 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 perfectly balance each other to keep average path length unchanged. The intuition behind the theorem 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 $k = \log(n) / \log(d)$ to reach all nodes, and given the exponential expansion, this would also be the average distance. The theorem shows that this is exactly how the average distance behaves even when the network is not a tree, even when we noise up the network so that nodes do not all have the same degree, and even when we add substantial homophily to the network. In proving this, there are two critical parts: first, the randomness of the nodes' degrees does not substantially alter the calculation (even if a power law distribution is admitted in expected degrees); and second, even though homophily may alter the structure of the network, the

¹⁰This result holds for a more general random network model, and is specialized to the multi-type random network model considered here for this statement.

shortest paths branching out from a given node are not much altered by homophily. The homophily affects which nodes are likely to be closer or further, but not the average distance.

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

The above results tell us that average distance is not affected by homophily, and diameter is affected only up to a fixed finite factor, provided there is some minimal level of inter-group connectivity. Thus average-path based communication processes are not affected by homophily but are affected by the connectivity of a society.

3.2 Markovian Processes: Linear Updating and the Random Walk

When we turn to the other forms of communication, there are substantial effects of homophily. It is not simply average distance that matters, but the *relative numbers* of paths between different nodes that matters.

To state the conceptual results most cleanly, we specialize to the islands model. The results extend to the more general multi-type random networks; those extensions require a somewhat longer exposition and so are discussed in the next section.

In the context of the islands model, let us define two measures of homophily. The (unnormalized) homophily is defined as

$$H = \frac{p_s}{p}$$

and captures how much more probable a link to a node of one's own type is compared to other types. This varies between 0 and m , the number of islands. If a node only links to same-type nodes, then the average linking probability p becomes p_s/m and so $H = m$, while if a node only links to nodes of other types, then $p_s = 0$ and so $H = 0$. We can also normalize the measure by dividing by the number of islands m ; the normalized homophily is thus defined as

$$h = \frac{p_s}{mp}.$$

Thus, h is the fraction of a node's links that are expected to be to agents of the same type.

If we index a sequence of societies by their cardinalities n , then the following theorem

summarizes the main conclusions of how homophily affects consensus and mixing times. The details and proof in a more general setting appear in the next section.

THEOREM 2. *In the equal-sized islands model, if $p(n)n/\log^2(n) \rightarrow \infty$ and $h(n)$ is bounded away from 1, then for any $\delta < 1$, high enough n ensures that the following is true with arbitrarily high probability:*

$$\frac{(1 - \delta) \log(n)}{2 \log(\frac{m-1}{H-1})} \leq \text{CT}(\gamma/n^2; \mathbf{A}(\mathbf{P}, \mathbf{n})) \leq \frac{(1 + \delta) \log(n)}{\log(\frac{m-1}{H-1})}$$

and

$$\frac{(1 - \delta) \log(n)}{\log(\frac{m-1}{H-1})} \leq \text{MT}(\gamma/n; \mathbf{A}(\mathbf{P}, \mathbf{n})) \leq \frac{\frac{3}{2}(1 + \delta) \log(n)}{\log(\frac{m-1}{H-1})}.$$

Theorem 2 provides us with a precise relationship between homophily and consensus and mixing times. As homophily increases, both consensus time and mixing time increase. In particular, both the consensus time $\text{CT}(\gamma/n^2; \mathbf{A})$ and the mixing time $\text{MT}(\gamma/n; \mathbf{A})$ are proportional (up to a fixed factor) to

$$\frac{\log(n)}{\log(\frac{m-1}{H-1})}.$$

This is true independently of connectivity and does not have any specific requirements about how many types (islands) there are. If m becomes large, then this further simplifies and both the consensus time $\text{CT}(\gamma/n^2; \mathbf{A})$ and the mixing time $\text{MT}(\gamma/n; \mathbf{A})$ are proportional (up to a fixed factor) to

$$\frac{\log(n)}{\log(1/h)}.$$

When homophily is low, then marginal increases in homophily has only a small effect. But as homophily grows large (H/m or h closer to 1), the magnitude of the marginal effect of increased homophily becomes very large.

3.3 Discussion

The interesting and intuitive contrast is the comparison between what matters for shortest-path and Markovian communication processes. Comparing Theorems 1 and 2, we see that, in the context of communication based on shortest paths, homophily has no effect while average connectivity is critical. In contrast, when considering Markovian processes, we see that homophily is critical while average connectivity is irrelevant.

This is an intuitive difference. With shortest-path communication, by definition, only the average distance between pairs of agents matters; it does not matter who is close and who is far from a given agent. Adding homophily changes network structure on the second dimension – introducing a certain kind of clustering (see Jackson (2008a) for more on clustering) – but it does not change the average lengths of the shortest paths branching out from each node. Agents just end up very close to other agents of their own types, who then provide connections to other types. While these last statements are not obvious, Theorem 1 asserts that the various effects interact in just the way needed to keep average distances the same.

With Markovian communication, on the other hand, the time to convergence is determined by how homophilous the network is, not by the density of its connections, provided that a certain (low) density threshold is met. A network in which the underlying links are formed without discrimination will not have any way to support long-term heterogeneity away from the steady state, even if the network is fairly sparse. On the other hand, a network with clearly defined islands of clustered agents will be able to maintain long-term differences in beliefs or behaviors. Each island will converge to its own metastable state and stay there for a long time, disagreeing with the other islands. Increasing the overall number of links while maintaining the homophily will not speed up convergence. The ratio of same-type connections to different-type connections for a given node will remain essentially the same, and so the network will be able to support the same disagreement based on the self-reinforcing effect within given types.

An interesting implication of this is the following. Consider a model where agents form links to others through a random search process, such as the one discussed in Currarini, Jackson and Pin (2009). Suppose that we consider a change in the matching technology – such as the introduction of social networking software – so that it becomes easier to search for agents of one’s own type. If agents have some preference for connecting to agents of their own types, this would lead to an increase in the overall density of links in the network, but would also cause the typical agent to spend a larger fraction of links on agents of the same type. What would the ultimate impact be in terms of communication? Our results imply that shortest-path based communication would become faster, but at the same time, Markovian processes of communication such as the linear updating process would converge more slowly!

4 Relating Communication Times, Second Eigenvalues, and Homophily

Theorem 2 is proven via a series of mathematical results which are of some interest in their own right. Again, as outlined in the introduction, the relation of speed of learning to homophily is established by breaking things into two parts: how speed relates to second eigenvalues, and how second eigenvalues relate to homophily. The key result that unlocks the second part of this puzzle is a representative agent theorem that shows that the second eigenvalue of a random network is, asymptotically, only dependent on the underlying probabilities of linking between different types of agents. The extra noise of which specific agents are linked to which others is essentially irrelevant in a large network. Only the broad patterns of linking across different types are important. Once we have proved this, we can relatively easily deduce results relating second eigenvalues to homophily, and then complete the picture relating speed to homophily.

The outline is summarized in Figure 1. A specific roadmap of the technical results is as follows.

Measures of convergence speed form the top layer of Figure 1 and summary statistics related to large-scale network structure form the bottom layer. The middle layer is the spectral intermediary that allows us to tie everything together. It is well-known that the second eigenvalue of a network's Markov matrix is a good proxy for the convergence speed of linear updating processes. Thus, in relating the top and middle layers of Figure 1, we use standard spectral results from the Markov chain literature to provide upper and lower bounds on mixing and consensus times. The results on mixing time, which is a central concept in Markov chain theory that captures how long it takes a process on a network to become random, are completely standard. We also define a notion of consensus time, which is essentially the time required for the mean squared deviation from consensus beliefs to get small. Bounding this quantity requires adapting standard results in a straightforward way (Lemma 3) but it turns out that in the multi-type random graph setting, tighter bounds than usual can be obtained for it (Proposition 6).

The main novel technical work of the paper concerns the relationship between the middle and bottom layers of the figure. First, we prove a general result, Theorem 3, which shows that in large multi-type random network, the study of the second eigenvalue of the entire network can be reduced to a computation based on a representative agent matrix which contains only one agent for each type. Building on this, we relate the second eigenvalue to more concrete measures of homophily. One that turns out to be particularly well suited to

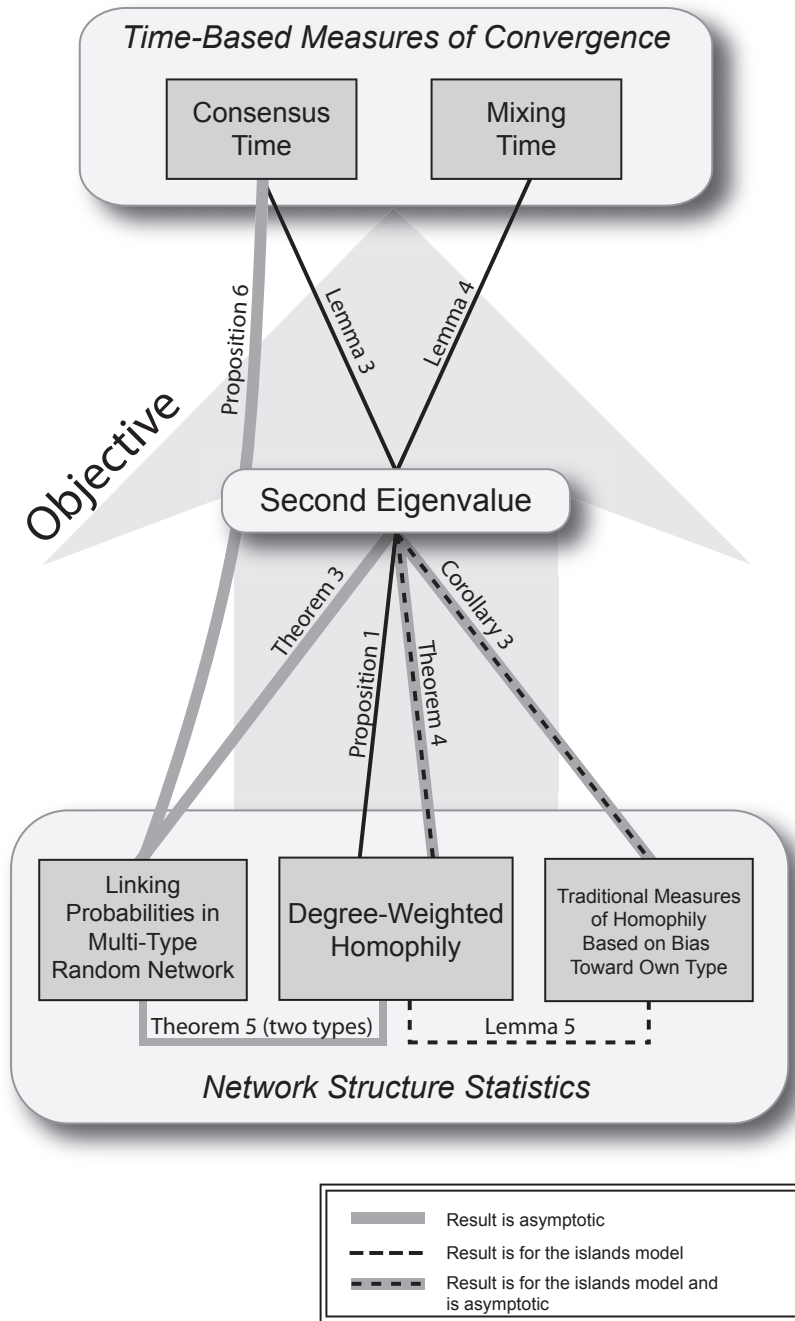


Figure 1: The paper’s conceptual structure. A line between two quantities indicates that one is used to bound or characterize the other; quantities located lower in the diagram are used to characterize the ones that are higher.

the study of eigenvalues, and hence of convergence, is a new quantity called degree-weighted homophily (DWH). This quantity measures the relative advantage of same-type links over different-type links, but does so in a way that takes into account different degrees and group sizes. Proposition 1 shows that, in arbitrary networks, this quantity always provides a lower bound on second eigenvalue, and hence consensus time. These results already entail that homophily slows learning and provide general tools for studying the relationship in arbitrary multi-type networks. For more concrete characterizations in an important special case, we turn to networks in which agents split into equally sized “islands”, where each island is a different type. Agents only discriminate based on whether someone else is inside or outside of their own islands. In this case, for large networks, we can exactly characterize second eigenvalues for large networks. These results can be stated in terms of a close relative of DWH (Theorem 2) and in terms of more traditional unweighted measures of homophily (Corollary 3).

Putting the transitions between the layers together, we end up with a tight relationship for large island networks between homophily and the speed of learning, which then leads to the conclusions described above. These are summarized in Theorem 2.

4.1 Relating Consensus and Mixing Times to Second Eigenvalues

We first state results that give fairly precise bounds on how the consensus time and mixing time of a matrix depends on its second eigenvalue. For any stochastic matrix \mathbf{T} , let $1 = \lambda_1(\mathbf{T}), \dots, \lambda_n(\mathbf{T})$ be its eigenvalues sorted by magnitude in decreasing order.

4.1.1 Consensus Time

LEMMA 3. *Assume \mathbf{A} is connected, let $\lambda_2(\mathbf{T}(\mathbf{A}))$ be the second largest eigenvalue in magnitude of $\mathbf{T}(\mathbf{A})$, and let \mathbf{s} be the (unique) steady-state distribution, with $\min_i s_i = \underline{s}$. If $\lambda_2(\mathbf{T}) \neq 0$, then for any $0 < \varepsilon \leq 1$:*

$$\left\lfloor \frac{\log(1/4\varepsilon) - \log(1/\underline{s})}{2 \log(1/|\lambda_2(\mathbf{T})|)} \right\rfloor \leq \text{CT}(\varepsilon; \mathbf{A}) \leq \left\lceil \frac{\log(1/\varepsilon)}{2 \log(1/|\lambda_2(\mathbf{T})|)} \right\rceil.$$

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 so we have a quite precise characterization in terms of the spectrum of the underlying social network.

The proof of this result follows fairly standard techniques from the spectral literature and the proof appears in the appendix.

The lower bound in Lemma 3 includes a term $\log(1/\underline{s})$ which can grow as n grows. We improve on the lower bound in Proposition 6 in the appendix, where we take advantage of the random graph structure to obtain a lower bound that does not depend on \underline{s} or n .

4.1.2 Mixing Time

Next, let us consider mixing time. Again, we can derive bounds based on the second eigenvalue. In this case, there is a difference that reflects the difference in the norms associated with these measures.

The following lemma is adapted from Montenegro and Tetali (2006) (see their Section 2.4: “Does Reversibility Matter”).

LEMMA 4. *[Montenegro and Tetali (2006)] Assume \mathbf{A} is connected. Let $\lambda_2(\mathbf{T})$ be the second largest eigenvalue in magnitude of \mathbf{T} , and \mathbf{s} be the (unique) steady-state distribution, with $\min_i s_i = \underline{s}$. If $\lambda_2(\mathbf{T}) \neq 0$, then for any $0 < \varepsilon \leq 1$:*

$$\frac{\log(\frac{1}{2\varepsilon})}{\log(\frac{1}{|\lambda_2(\mathbf{T})|})} \leq \text{MT}(\varepsilon; \mathbf{A}) \leq \frac{\log(\frac{1}{2\varepsilon}) + \log(1/\underline{s})/2}{\log(\frac{1}{|\lambda_2(\mathbf{T})|})}.$$

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

4.2 Relating Second Eigenvalues to Network Structure

4.2.1 A Representative Agent Theorem

We now present our main technical result, a “representative-agent” theorem that allows us to analyze the convergence of a multi-type random graph by studying a much smaller graph in which there is only one node for each *type* of agent. We show that under some conditions on the minimum expected degree, the second eigenvalue of most any realized multi-type random graph converges in probability to the second eigenvalue of this representative-agent matrix. This result is useful for dramatically simplifying computations of approximate consensus times and mixing times, both in theoretical results and in empirical settings, as now the random second eigenvalue 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.

Recalling the notation from Section 2.2, let $d_{k\ell}(\mathbf{P}, \mathbf{n}) = n_\ell P_{k\ell}$ the expected number of links that a node of type k will have with nodes of type ℓ and let $d_k(\mathbf{P}, \mathbf{n}) = \sum_\ell d_{k\ell}(\mathbf{P}, \mathbf{n})$

be the expected degree of a node of type k . Let $\mathbf{Q}(\mathbf{P}, \mathbf{n})$ be a matrix of the same dimensions as \mathbf{P} with entries

$$Q_{k\ell} = \frac{d_{k\ell}(\mathbf{P}, \mathbf{n})}{d_k(\mathbf{P}, \mathbf{n})}.$$

So, $Q_{k\ell}$ is the expected relative fraction of links that a node 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 rather than individual nodes, and second, it works with expected fractions rather than realized values.¹¹

THEOREM 3. *Consider a multi-type random network described by (\mathbf{P}, \mathbf{n}) . For any $\delta > 0$ there exists K such that if $\min_k d_k(\mathbf{P}, \mathbf{n}) > K \log^2 n$ then*

$$|\lambda_2(\mathbf{T}(\mathbf{A}(\mathbf{P}, \mathbf{n}))) - \lambda_2(\mathbf{Q}(\mathbf{P}, \mathbf{n}))| \leq \delta, \tag{1}$$

with probability at least $1 - \delta$.

Theorem 3 is a law of large numbers for spectra of multi-type random graphs. Such 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 precise global structure. The closest antecedent to this particular theorem is by Chung, Lu and Vu (2004) for networks without homophily. This theorem is the first of its kind to apply to a model that allows homophily and the associated heterogeneities in linking probabilities. We employ a similar strategy of proof, which relies on decomposing the random matrix representing our graph into two pieces: an “orderly” piece whose entries are given by linking 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

¹¹There is one technical issue which does not substantially affect any results but which deserves some comment. If one is thinking of a multi-type random network as describing the relationships relevant for a boundedly rational process of belief updating, then the self-weights assumed in our formulation of the model are unnatural. In particular, if we think of a link from i to j as capturing the fact that i has access to the belief of j , then all nodes should have self-links; in the model as we have described it, a node of type k has a self link with probability only P_{kk} . It turns out that formulating the model so that all nodes always have self-links does not change any of the results presented in this section, except for a small change in asymptotic rates of convergence in Theorem 3 and its consequences. This is because the spectral norm of the difference between the matrix we work with in the proof and the matrix with self-links decays to 0, assuming that minimum degree is unbounded. The same is true if we forbid self-links, which is natural in the dialect-choice game of Section 2.3.2 or the Facebook random walk of Section 2.3.3. In short, the results are not sensitive to how we model self links, provided that everyone treats all neighbors, including oneself, equally.

note that computing the second eigenvalue of the orderly part requires dealing only with a representative-agent matrix, which will usually be small.

The usefulness of Theorem 3 becomes evident through a series of its implications. First, it can be used to tighten the lower bound in Lemma 3, so that second eigenvalues become an even better proxy for consensus time; this is done in Proposition 6 in the appendix. In the next section, we also use the theorem to derive expressions allowing us to understand how homophily affects consensus and mixing time in the islands model. Lastly, we use it to show that the degree weighted homophily bound in Proposition 1 is tight.

4.2.2 Non-Spectral Measures of Homophily

The relationships between consensus and mixing time relate to the second eigenvalue of the updating matrix, and through our representative agent theorem, to the second eigenvalue of the matrix of probabilities of connection across types in the random networks model. These characterizations are still somewhat abstract as the second eigenvalue is an implicitly defined statistic that may be hard to grasp. In order to understand the implications of the structural feature of homophily on mixing and consensus time, we need to develop some understanding of how homophily affects the second eigenvalue.

We begin with some general definitions of homophily, and later specialize to multi-type random networks.

General Networks and Degree-Weighted Homophily Let us partition N into two subsets, M and M^c . First, we define a notion of the weight between two groups.

DEFINITION 3. *Given $\mathbf{T} = \mathbf{T}(\mathbf{A})$ and two subsets of nodes, $B, C \subseteq N$, let*

$$W_{B,C} = \frac{\sum_{\substack{i \in B \\ j \in C}} T_{ij} T_{ji}}{|B||C|}.$$

$W_{B,C}$ keeps track of the relative weight between two sets of nodes B and C , and is a measure that ranges between 0 and 1. The weight of an edge is proportional to the reciprocal of the product of the degrees of the nodes on its ends: when an edge is between two nodes that have many neighbors, it doesn't count for much, but when it is between two that have few neighbors, it counts for a lot. The weight also depends on group size: individual edges within larger groups matter less than those within smaller groups. With this definition in hand, we define a notion of degree-weighted homophily.

DEFINITION 4. Given any $\emptyset \subsetneq M \subsetneq N$, let the degree-weighted homophily of the network \mathbf{A} relative to M be defined by

$$\text{DWH}(M; \mathbf{A}) = \frac{W_{M,M} + W_{M^c,M^c} - 2W_{M,M^c}}{\frac{1}{|M|^2} \sum_{i \in M} \frac{1}{d_i(\mathbf{A})} + \frac{1}{|M^c|^2} \sum_{i \in M^c} \frac{1}{d_i(\mathbf{A})}}, \quad (2)$$

where the W 's are relative to $\mathbf{T}(\mathbf{A})$.

The term in the numerator keeps track of how much of the weight in \mathbf{T} falls within M and within M^c , and how much weight goes between these sets of nodes. So, links within the group M or its complement M^c increase the degree weighted homophily and links between the two groups decrease it. The term in the denominator is a normalizing value which guarantees¹² that this quantity is always between -1 and 1 .

To see that the degree-weighted homophily has an intuitive interpretation, consider a very simple special case. Suppose $|M| = n/2$ and \mathbf{A} corresponds to a regular graph, where all degrees are equal. Then

$$\text{DWH}(M; \mathbf{A}) = \frac{\#(\text{within-group edges}) - \#(\text{between-group edges})}{\#(\text{total edges})}. \quad (3)$$

The theoretical justification for the usefulness of this measure is that it provides a lower bound on the magnitude of the second eigenvalue of \mathbf{T} , and in the limit a tight bound.

Let

$$\text{DWH}(\mathbf{A}) = \max_{\emptyset \subsetneq M \subsetneq N} |\text{DWH}(M; \mathbf{A})|.$$

Thus, the degree weighted homophily of a given network is the maximum level of degree homophily across different possible splits of the network.¹³

PROPOSITION 1. Assume that \mathbf{A} is connected. Then

$$|\lambda_2(\mathbf{T}(\mathbf{A}))| \geq |\text{DWH}(\mathbf{A})| \quad (4)$$

Combining this with Lemmas 3 and 4 we see that degree weighted homophily provides a lower bound on the consensus and mixing times.

¹²This can be verified by using the expression of DWH as a quadratic form in the proof of Proposition 1 below and then noting that the spectral norm of the matrix $\mathbf{T}(\mathbf{A})$ is 1.

¹³This has intuitive relationships to a weighted version of a min cut, although this degree weighted homophily measure turns out to be the right one for our purposes.

COROLLARY 2. *Assume \mathbf{A} is connected. Then for any $0 < \varepsilon < 1$,*

$$\text{CT}(\varepsilon; \mathbf{A}) \geq \left\lfloor \frac{\log(1/(4\varepsilon)) - \log(1/\underline{s})}{\log(1/\text{DWH}(\mathbf{A}))} \right\rfloor,$$

and

$$\text{MT}(\varepsilon; \mathbf{A}) \geq \left\lfloor \frac{\log(1/(2\varepsilon))}{\log(1/\text{DWH}(\mathbf{A}))} \right\rfloor.$$

Compared with Lemmas 3 and 4, the results have only one inequality each. This is because DWH only provides a lower bound on eigenvalues, and hence on the convergence times. In the next subsection, we supply asymptotic upper bounds in the islands setting.

The Islands Model and Simpler Measures of Homophily In order to develop the clearest and most intuitive relationships between homophily and the resulting consensus and mixing times, we now examine a specific case of the multi-type random network model. Recall the islands model from Section 2.2.1, and consider the case where there are $m \geq 2$ equally-sized groups. In all the results of this section, we will consider only n divisible by m , and the results will concern limits as $n \rightarrow \infty$. All quantities (probabilities, homophilies, etc.) are implicitly indexed by n , but we suppress this indexing unless it is important to emphasize it. Let p_s and p_d be the probability of links within and across types, respectively, and p be the overall probability of links.

Let $\text{EDWH}(m, p_s, p_d)$ denote the expected degree weighted homophily in the islands model where we calculate this relative to the *expected* number of links within and across islands. That is, if we have a collection of k islands, M , let

$$EW_{M,M} = \frac{[p_s k + p_d k(k-1)]/d^2}{k^2}$$

and

$$EW_{M,M^c} = \frac{p_d k(m-k)/d^2}{k(m-k)} = p_d/d^2,$$

where $d = pn$ is the expected degree, and EW_{MM} and EW_{MM^c} are the expected versions of W_{MM} and W_{MM^c} . Then, we have an expected variation of degree weighted homophily:

$$\text{EDWH}(M; m, p_s, p_d) = \frac{EW_{M,M} + EW_{M^c,M^c} - 2EW_{M,M^c}}{\frac{1}{|M|^2} \sum_{i \in M} \frac{1}{d} + \frac{1}{|M^c|^2} \sum_{i \in M^c} \frac{1}{d}}.$$

Let $I(n)$ denote the subsets of nodes that are collections of islands, so that if $M \in I(n)$ then

any node in M is of a different type from any node in M^c . Let

$$\text{EDWH}(m, p_s, p_d) = \max_{M \in I(n)} \text{EDWH}(M; m, p_s, p_d).$$

This is not quite the expected degree weighted homophily, since we are working with expectations in the numerator and denominator.

We can now prove the following theorem which lets us deduce how consensus and mixing time depend on degree weighted homophily in the islands (see Corollary 4), since it establishes how the second eigenvalue relates to homophily and we have already established how consensus and mixing time relate to the second eigenvalue.

THEOREM 4. *In the equal-sized islands model, if $p(n)n/\log^2(n) \rightarrow \infty$ then*

$$|\lambda_2(\mathbf{T}(\mathbf{A}(\mathbf{P}, \mathbf{n}))) - \text{EDWH}(m, p_s, p_d)| \xrightarrow{p} 0. \quad (5)$$

Theorem 4 provides us with limiting expressions for the second eigenvalue as a function of homophily.

Let

$$H = \frac{p_s}{p}$$

capture how much more probable a link to own type is compared to other types, and

$$h = \frac{p_s}{mp}$$

be the relative fraction of links to own type. If we index a sequence of societies by their cardinalities n , then the following lemma establishes the relation between degree weighted homophily and the relative probabilities of links and the number of islands.

LEMMA 5. *In the islands model with $m \geq 2$ equal-sized groups and probabilities of links within and across types p_s and p_d , respectively, the degree weighted homophily is*

$$\text{EDWH}(m, p_s, p_d) = \frac{p_s - p_d}{p_s + (m-1)p_d} = \frac{H - 1}{m - 1}.$$

Moreover,

$$\text{EDWH}(m, p_s, p_d) = \text{EDWH}(M; m, p_s, p_d)$$

for all $M \in I(n)$, so the grouping of the islands is irrelevant in calculating the homophily. If the number of islands $m(n)$ diverges then $|\text{EDWH}(m, p_s, p_d) - h| \rightarrow 0$.

From Theorem 4 and Lemma 5, the following corollary, characterizing second eigenvalues in terms of traditional measures of homophily in the islands model, follows immediately.

COROLLARY 3. *In the islands model with $m \geq 2$ equal-sized groups, if $p(n)n/\log^2(n) \rightarrow \infty$ and probabilities of links within and across types p_s and p_d , respectively,*

$$\left| \lambda_2(\mathbf{T}(\mathbf{A}(\mathbf{P}, \mathbf{n}))) - \frac{H-1}{m-1} \right| \xrightarrow{p} 0. \quad (6)$$

If the number of islands diverges, then $|\lambda_2(\mathbf{T}(\mathbf{A}(\mathbf{P}, \mathbf{n}))) - h| \rightarrow 0$.

Combining Lemmas 3 and 4 with Corollary 3 we have the following summary of the asymptotic behavior of consensus and mixing time in the islands model, which is Theorem 2.

COROLLARY 4. *In the equal-sized islands model, if $p(n)n/\log^2(n) \rightarrow \infty$ and $h(n)$ is bounded away from 1, then for any $\delta < 1$, with a probability going to 1:*

$$\frac{(1-\delta)\log(n)}{2\log\left(\frac{m-1}{H-1}\right)} \leq \text{CT}(\gamma/n^2; \mathbf{A}(\mathbf{P}, \mathbf{n})) \leq \frac{(1+\delta)\log(n)}{\log\left(\frac{m-1}{H-1}\right)}$$

and

$$\frac{(1-\delta)\log(n)}{\log\left(\frac{m-1}{H-1}\right)} \leq \text{MT}(\gamma/n; \mathbf{A}(\mathbf{P}, \mathbf{n})) \leq \frac{\frac{3}{2}(1+\delta)\log(n)}{\log\left(\frac{m-1}{H-1}\right)}.$$

The condition that h is bounded away from 1 rules out the case that all but a vanishing fraction of links are within islands. If that is the case, then the islands can become disconnected with a nontrivial probability and the mixing and consensus times diverge. The consideration of setting $\varepsilon = 1/n^2$ for consensus time and $\varepsilon = 1/n$ for mixing time ensure that the convergence is within the order of magnitude of the weight on any given node. We can rewrite these expressions as

$$\frac{(1-\delta)\log(n)}{2|\log(\text{EDWH}(m, p_s, p_d))|} \leq \text{CT}(\gamma/n^2; \mathbf{A}(\mathbf{P}, \mathbf{n})) \leq \frac{(1+\delta)\log(n)}{|\log(\text{EDWH}(m, p_s, p_d))|}$$

and

$$\frac{(1-\delta)\log(n)}{|\log(\text{EDWH}(m, p_s, p_d))|} \leq \text{MT}(\gamma/n; \mathbf{A}(\mathbf{P}, \mathbf{n})) \leq \frac{\frac{3}{2}(1+\delta)\log(n)}{|\log(\text{EDWH}(m, p_s, p_d))|}.$$

Corollary 4 provides us with a fairly precise relationship between homophily and consensus and mixing times, as further discussed in previous sections.

The above results presume equal sized islands. We also provide a result for unequal sizes for the case of two islands. This shows that the degree weighted homophily bound in Proposition 1 is tight.

THEOREM 5. *Suppose $\mathbf{n} = (\lceil f_1 n \rceil, \lfloor (1 - f_1)n \rfloor)$ with $0 < f_1 < 1$ and*

$$\mathbf{P} = \begin{bmatrix} p_s & p_d \\ p_d & p_s \end{bmatrix},$$

where all the entries of this matrix are positive. Then

$$\text{plim}_{n \rightarrow \infty} \lambda_2(\mathbf{T}(\mathbf{A}(\mathbf{P}, \mathbf{n}))) = \frac{p_s}{p_s + p_d(f_1^{-1} - 1)} - \frac{p_d}{p_s(f_1^{-1} - 1) + p_d} = \text{plim}_{n \rightarrow \infty} \text{DWH}(N_1; \mathbf{A}(\mathbf{P}, \mathbf{n})),$$

where N_1 denotes the first $\lceil f_1 n \rceil$ nodes.

5 Consensus and Mixing Times in the Adolescent Health Data

In this section, we examine consensus time and mixing times in 84 social networks from the Adolescent Health dataset¹⁴ and show how the patterns in that data illustrate our conclusions. For each of the 84 schools, the dataset includes information on each student's grade, gender and race. In addition, each student was asked to name his or her closest male and female friends.¹⁵ Using the reported friendship networks (linking two individuals if either named the other as friend) we compute consensus and mixing times. We can also examine traditional and degree weighted homophily measures based on the observed characteristics: grade, race and sex. Grade is the year in school, and ranges from 6 to 12, as most of the schools include 6 different years of students. Race is self-reported as Asian, black, Hispanic, white, or other (and these were the only categories permitted). Sex is self-reported as male

¹⁴Add Health is a program project designed by J. Richard Udry, Peter S. Bearman, and Kathleen Mullan Harris, and funded by a grant P01-HD31921 from the National Institute of Child Health and Human Development, with cooperative funding from 17 other agencies. Persons interested in obtaining data files from Add Health should contact Add Health, Carolina Population Center, 123 W. Franklin Street, Chapel Hill, NC 27516-2524 (addhealth@unc.edu). We thank James Moody for making available the data organized in Pajek files for the 84 schools.

¹⁵The number of friends reported was capped at five of each type, or ten in total. Less than ten percent of the students hit the caps, but that still censors the data. This design feature makes homophilies computed based on gender somewhat less reliable than the others, since it would tend to equalize the numbers of reported male and female friends, even if there were strong homophily present.

or female.

It is worth commenting briefly on the nature of the illustration that our empirical computations provide for the theoretical work. There are actually three aspects to the full spectrum of our results that can be tested.

- First, there is a question of whether substantial information about the impact of homophily can be captured by examining relatively simple definitions of types.
- Second, many of our results are asymptotic and there is a question about whether our bounds on how consensus time and mixing time relate to homophily will be useful in finite samples of medium size.
- Third, there is a question of whether or not people actually communicate in ways that are captured by the learning model and random Markov models that underlie consensus and mixing time.

Our empirical analysis answers the first two questions in the affirmative. In particular, the multi-type random network model is a good fit for these social networks when it comes to investigating consensus and mixing times, and gets a good deal of explanatory power from very basic definitions of types. Our main claim — that the study of the convergence of Markovian processes on large networks can be reduced to simple computations about homophily — is not merely an asymptotic, theoretical claim, but one that holds up well when applied to the data. To see this more clearly, consider some of the typical building blocks that went into establishing the relationship between convergence and homophily: Lemma 3, Theorem 3, and Corollary 4, for example. Each result either provides inequalities or statements about asymptotic convergence. *A priori*, the data might be badly behaved with respect to either of these. It might stay within the inequalities in a noisy way (oscillating randomly between the bounds). It might also take very large networks for the asymptotic results to kick in. Lastly, it might even be the case that the multi-type random graph model captures none of the salient structure of these social networks. If any of these happened, then there might be very weak, nonexistent, or “wrong way” correlations in the quantities that we studied empirically in this section. The fact that the correlations are quite strong and correspond to our predictions shows that the relationships suggested by the inequalities and asymptotic theory are relevant. In particular, these high school friendship networks seem to have many of the salient features of well-behaved multi-type random graphs, and much can be captured with simple definitions of types.

Whether or not these models of updating and communication shed light on actual social behavior — that is, on how people actually communicate in or navigate networks — is obviously

an important question, but one that requires additional (longitudinal) data and is left for future investigations.

The first step is to test the relationship summarized in Theorem 2 between convergence and homophily in the equal-sized islands model. We begin the building blocks of the theorem to rearrange it into an inverted form that compensates for the extreme behavior of the convergence times at high homophilies and makes the quantities amenable to linear regressions. In particular, we define

$$\rho(X) = \exp\left(-\frac{\log(n)}{X}\right).$$

By Lemmas 3 and 4, $\rho(\text{CT}(\gamma/n^2; \mathbf{A}))$ and $\rho(\text{MT}(\gamma/n; \mathbf{A}))$ are approximately the second eigenvalue of $\mathbf{T}(\mathbf{A})$, which, by Corollary 3, is well approximated by $\frac{H-1}{m-1}$. For instance, if X is an empirical measurement of a consensus time for some choice of $\varepsilon = \gamma/n^2$, then $\rho(X)$ can be thought of as an imputed per-step rate of convergence. Thus, we run regressions of $\rho(\text{CT}(\gamma/n^2; \mathbf{A}))$ and $\rho(\text{MT}(\gamma/n; \mathbf{A}))$ on $\frac{H-1}{m-1}$.

The regressions include an intercept term. This is not in Theorem 2 or its constituent parts. However, it turns out that if there is additional homophily within each island on dimensions not reported in the data, then there will be an intercept term in the model. Details are in Section 7.3.1.

Lastly, we removed two data points whose consensus and mixing times exceeded our algorithms' capacity. These networks (schools number 53 and 57) had very large consensus times (on the order of several thousand), so computing them precisely was infeasible. These would not substantially change the results of the first two regressions, since for those purposes, a consensus or mixing time of about 1000 is essentially infinite. So, from now on, we work with the 82 data points excluding those schools. The results are presented in Tables 1 and 2, as well as Figure 2.

We begin with the finest definition of type available in the data. Thus, we consider a "type" to be a specific combination of race, grade, and sex: so for instance a type would be all female Asians in grade 9. Thus, in a high school with two sexes, four races, and four grades there are thirty two types.

Table 1: Dependent variable = $\rho(\text{CT}(0.1/n^2; \mathbf{A}))$
 ($N = 82$)

Variable	Coefficient (<i>t</i>-statistic)
Intercept	0.870 (61.2)
$\frac{H-1}{m-1}$ for “type”	0.297 (4.91)
R^2	0.231

Table 2: Dependent variable = $\rho(\text{CT}(0.1/n^2; \mathbf{A}))$
 ($N = 82$)

Variable	Coefficient (<i>t</i>-statistic)
Intercept	0.861 (66.9)
$\frac{H-1}{m-1}$ for “type”	0.287 (5.23)
R^2	0.255




Figure 2: The relationship between $\rho(\text{CT}(0.1/n^2; \mathbf{A}))$ and $\frac{H-1}{m-1}$ computed based on the finest-grained type data available (i.e. a type is a race-grade-sex tuple).

The R^2 in the above regressions shows that the homophily among these types accounts for roughly a quarter of the variation in consensus and mixing times in the data. This is quite high in view of the fact that many qualities that determine network formation – such as interests, extracurricular activities, etc. – are not captured by these data.

We also explore how much of the variation can be explained by the even simpler definitions of types. For example, for grades, out of the three characteristics captured in the data, has the greatest variation in homophily. The grades also have approximately equal sizes in most of the schools, so that it is legitimate to use the formulas from the equal-sized islands model. The results are then reported in Tables 3 and 4, as well as Figure 3.

Table 3: Dependent variable = $\rho(\text{CT}(0.1/n^2; \mathbf{A}))$
($N = 82$)

Variable	Coefficient (<i>t</i>-statistic)
Intercept	0.809 (32.6)
$\frac{H-1}{m-1}$ for grade	0.209 (5.21)
R^2	0.253

Table 4: Dependent variable = $\rho(\text{CT}(0.1/n^2; \mathbf{A}))$
($N = 82$)

Variable	Coefficient (<i>t</i>-statistic)
Intercept	0.825 (34.6)
$\frac{H-1}{m-1}$ for grade	0.163 (4.24)
R^2	0.184

The fit is similar in quality (though weaker in the case of mixing time) to that obtained from the finest definitions of type.

We can also examine some other implications of the theory. By focusing on Theorem 4 instead of Theorem 2 we can replace $\frac{H-1}{m-1}$ in the above regressions by DWH. The increase in explanatory power comes from the fact that DWH does not presume equal-sized groups, and




Figure 3: The relationship between $\rho(\text{CT}(0.1/n^2; \mathbf{A}))$ and $\frac{H-1}{m-1}$ for grade.

thus captures the fact that different islands may be of different sizes. We compute DWH based on the three different observed characteristics in the data and use them to estimate $\rho(\text{CT}(\gamma/n^2; \mathbf{A}))$ and $\rho(\text{MT}(\gamma/n; \mathbf{A}))$. For a given definition of type, we take the DWH over all nontrivial partitions that never separate two agents of the same type. For example, $\text{DWH}_{\text{grade}}(\mathbf{A})$ is the DWH taken over all partitions which have some grades on one side and the rest of the grades on the other. In the tables below, we refer to this quantity as Grade DWH, and similarly with other type classifications. More generally, given some definition θ of type such that $\theta : N \rightarrow C$ maps agents to types, we define

$$\text{DWH}_{\theta}(\mathbf{A}) := \max_{\substack{\theta^{-1}(M) \cap \theta^{-1}(M^c) = \emptyset \\ \emptyset \subsetneq M \subsetneq N}} \text{DWH}(M; \mathbf{A}).$$

For this analysis, we do not compute DWH for types, as with thirty or so different groups, the number of different partitions is such that the computations become infeasible.

Here we are measuring realized degree weighted homophily, DWH, rather than its “expected” analogue, EDWH. This is because EDWH is not available to us, so, as usual, we replace it by the sample analogue. As shown by the following lemma, this is valid asymptotically.

LEMMA 6. *Consider the islands model with $m(n) \geq 2$ equal-sized groups where $m(n)/n \rightarrow 0$, and probabilities of links within and across types p_s and p_d , respectively, and consider any sequence of groupings of islands $M(n) \in I(n)$. Then*

$$|\text{EDWH}(M(n), m, p_s, p_d) - \text{DWH}(M(n), \mathbf{A}(\mathbf{P}, \mathbf{n}))| \xrightarrow{p} 0,$$

and so

$$|\text{EDWH}(m, p_s, p_d) - \text{DWH}(M(n), \mathbf{A}(\mathbf{P}, \mathbf{n}))| \xrightarrow{p} 0.$$

Regressions of convergence rates on the DWH for the 82 networks are reported in Tables 5 and 6. Here, we run the regressions with an intercept term, which is motivated by the same idea as the one formally justifying the inclusion of an intercept in the first regressions of this section. We have not worked out the details formally. Moreover, one of our regressions includes three DWH explanatory variables – one for each dimension. Such an additively separable form is not justified by the theory but seems to track the data quite closely, so we include it here to point out a potential relationship which may be fruitful to examine further.

Table 5: Dependent variable = $\rho(\text{CT}(0.1/n^2; \mathbf{A}))$
 ($N = 82$)

Variable	Coefficient (<i>t</i> -statistic)			
	<i>All Homophilies</i>	<i>Grade Only</i>	<i>Gender Only</i>	<i>Race Only</i>
Intercept	0.644 (20.7)	0.716 (21.52)	0.886 (56.4)	0.916 (88.1)
Grade DWH	0.347 (7.75)	0.330 (6.64)	—	—
Gender DWH	0.137 (2.64)	—	0.231 (3.36)	—
Race DWH	0.105 (5.04)	—	—	.0663 (2.29)
R^2	0.545	0.356	0.123	0.0617

Consider Table 5. The first regression includes all the homophilies acting as independent variables, and the others have each type of homophily used as an explanatory variable on its own. The table shows two main things. First, all three homophilies are significant at the 2% level when the regression is run with all three explanatory variables. Second, grade homophily is doing most of the work in explaining the variation in convergence times; other kinds of homophily have a significant effect, but 36% of the variation can be explained by ignoring all but the grade information. This is illustrated in Figure 4, where we plot $\rho(\text{CT}(0.1/n^2; \mathbf{A}))$ versus grade homophily and draw the least-squares trend line corresponding to the second column of Table 5.




Figure 4: Rate of convergence of consensus time for the 82 friendship networks plotted against the degree weighted homophily in each of the networks calculated relative to grade.

Table 6: Dependent variable = Imputed convergence rate, $\rho(\text{MT}(0.1/n; \mathbf{A}))$
($N = 82$)

Variable	Coefficient (<i>t</i> -statistic)			
	<i>All Homophilies</i>	<i>Grade Only</i>	<i>Gender Only</i>	<i>Race Only</i>
Intercept	0.665 (22.4)	0.735 (23.3)	0.884 (60.2)	0.902 (95.6)
Grade DWH	0.309 (7.26)	0.284 (6.05)	—	—
Gender DWH	0.100 (2.02)	—	0.183 (2.84)	—
Race DWH	0.104 (5.23)	—	—	0.0696 (2.65)
R^2	0.511	0.314	0.0917	0.0807




Figure 5: Rate of convergence of mixing time for the 82 friendship networks plotted against the degree weighted homophily in each of the networks calculated relative to grade.

We repeat the analysis with mixing times, and the results turn out to be qualitatively the same, with slightly weaker values of R^2 , as summarized in Table 6. The corresponding plot for the rate of convergence of mixing time versus DWH based on grade is in figure 5.

The fact that consensus time and mixing time show similar results is not surprising, given that they both involve measures of distance between $\mathbf{T}(\mathbf{A})$ and its limit. To see this directly, we note the tight relationship between the two in Figure 6.




Figure 6: The relationship between consensus and mixing times.

While the above analysis has examined how imputed rates of convergence are affected by homophily, we could work with the consensus times and mixing times directly and compare them to the prediction of Theorem 2, which states that they should be approximately proportional to $\log(n)/\log(\frac{m-1}{H-1})$. When we perform such an analysis, we find results consistent with the theory, as pictured in Figures 8 and 7, where the slope coefficients are both significant at levels well below .001 and the intercept is constrained to be 0. However, the values of R^2 in these regressions are low (0.06 for CT and 0.01 for MT) because some extreme data points contribute very large error under this parameterization. It was for this reason that we changed the axes in the above regressions; the rescaling makes errors comparable across data points.




Figure 7: The relationship between consensus time and the prediction of Theorem 2 for grade.




Figure 8: The relationship between mixing time and the prediction of Theorem 2 for grade..

6 Concluding Remarks

Our results are built in several parts:

- (i) we relate communication processes to second eigenvalues largely building on standard spectral theory,
- (ii) we provide novel results relating second eigenvalues to homophily,
- (iii) we provide novel results relating homophily to random graphs, and
- (iv) finally, combining these results enables us to relate communication to homophily in random networks.

Our results show that homophily can substantially affect communication processes but that this depends both on the level and type of homophily and the type of communication. Communication based on shortest paths is essentially unaffected by homophily, while random walks and updating by averaging are affected in a well-identified and nonlinear manner. The underlying reason is that homophily does not change shortest paths, but affects the relative numbers of paths between nodes of different types. Interestingly, there is a complete reversal in the manner in which communication depends on the structure of connections:

- The speed of shortest path communication depends on link density but not homophily.
- The speed of Markovian processes (weighted averaging and random walks) depends on homophily but not link density.

We have also examined a set of 82 networks to see how the communication processes would operate on these networks, and how that relates to the observed homophily of the networks. The results show significant relationships that are as predicted by the theory, with increased homophily leading to increased consensus and mixing times according to the predicted formulas.

Our results suggest the importance of understanding homophily in order to understand communication and the functioning of a society. This is, of course, a first step and suggests many avenues for further research including endogenizing homophily and network structure, considering other sorts of communication, and examining other data.

References

- CHUNG, F. AND L. LU (2002): “The average distance in random graphs with given expected degrees,” *Proceedings of National Academy of Science*, 99, 15879–15882.
- CHUNG, F., L. LU, AND V. VU (2004): “The Spectra of Random Graphs with Given Expected Degrees,” *Internet Mathematics*, 1, 257–275.
- COPIC, J., M. O. JACKSON, AND A. KIRMAN (2005): “Identifying Community Structures from Network Data,” Revised September, 2008. Preprint, available at <http://www.stanford.edu/~jacksonm/netcommunity.pdf>
- CURRARINI, S., M. O. JACKSON, AND P. PIN (2009): “An Economic Model of Friendship: Homophily, Minorities and Segregation,” *Econometrica*, forthcoming.
- DEGROOT, M. H. (1974): “Reaching a Consensus,” *Journal of the American Statistical Association*, 69, 118–121.
- DEMARZO, P., D. VAYANOS, AND J. ZWIEBEL (2003): “Persuasion Bias, Social Influence, and Uni-Dimensional Opinions,” *Quarterly Journal of Economics*, 118, 909–968.
- DIACONIS, P. AND D. STROOCK (1991): “Geometric Bounds for Eigenvalues of Markov Chains,” *The Annals of Applied Probability*, 1, 36–61.
- FRENCH, J. (1956): “A Formal Theory of Social Power,” *Psychological Review*, 63, 181–194.
- GOLUB, B. AND M. O. JACKSON (2007): “Naïve Learning in Social Networks: Convergence, Influence, and the Wisdom of Crowds,” Preprint, available at <http://www.stanford.edu/~jacksonm/naivelearning.pdf>
- HARARY, F. (1959): “A Criterion for Unanimity in French’s Theory of Social Power,” in *Studies in Social Power*, ed. by D. Cartwright, Ann Arbor: University of Michigan Press.
- HORN, R. A. AND C. R. JOHNSON (1985): *Matrix Analysis*, Cambridge: Cambridge University Press.
- JACKSON, M. O. (2008a): “Average Distance, Diameter, and Clustering in Social Networks with Homophily,” in *Proceedings of the Workshop on Internet and Network Economics (WINE 2008)*, ed. by C. Papadimitriou and S. Zhang, Springer-Verlag, forthcoming.
- (2008b): *Social and Economic Networks*, Princeton, N.J.: Princeton University Press.

LAZARFELD, P. AND R. MERTON (1954): “Friendship as a social process: a substantive and methodological analysis,” in *Freedom and Control in Modern Society*, ed. by M. Berger, New York: Van Nostrand.

MCPHERSON, M., L. SMITH-LOVIN, AND J. COOK (2001): “Birds of a Feather: Homophily in Social Networks,” *Annual Review of Sociology*, 27, 415–44.

MONTENEGRO, R. AND P. TETALI (2006): “Mathematical Aspects of Mixing Times in Markov Chains,” *Foundations and Trends in Theoretical Computer Science*, 1, 237–354.

PLATO (1969): *The Republic: Books I-V, with an English Translation by Paul Shorey*, Cambridge, MA: Harvard University Press.

7 Appendix: Proofs

7.1 Background on Reversible Markov Chains

For completeness and to fix notation, we review very well-known results about Markov chains and self-adjoint matrices which form the foundation for our measures of convergence to consensus and bounds on the time required to converge. None of the material in this section is original; further background and references on these techniques can be found in Diaconis and Stroock (1991).

Symmetry or self-adjointness is often a useful property to have when working with eigenvalues and other spectral quantities of a matrix. While $\mathbf{T}(\mathbf{A})$ generally will not be symmetric, we can make it into a self-adjoint operator under a well-chosen inner product, as in Diaconis and Stroock (1991). For this we need a few definitions.

Given a probability distribution π on N , define

$$\langle \mathbf{v}, \mathbf{w} \rangle_{\pi} = \sum_i v_i w_i \pi_i.$$

This is just the Euclidean inner product weighted by the entries of the distribution.

DEFINITION 5. *A stochastic matrix \mathbf{T} satisfies detailed balance (equivalently, is reversible) relative to the distribution π over the nodes if, for every $i, j \in N$ we have*

$$\pi_i T_{ij} = \pi_j T_{ji}.$$

PROPOSITION 2. *If \mathbf{T} satisfies detailed balance relative to $\boldsymbol{\pi}$ then $\boldsymbol{\pi}$ is a stationary distribution for \mathbf{T} .*

Proof of Proposition 2: Observe that

$$\sum_j T_{ji}\pi_j = \sum_j T_{ij}\pi_i = \pi_i \sum_j T_{ij} = \pi_i,$$

where the first equality uses the definition of detailed balance. ■

PROPOSITION 3. *The stochastic matrix \mathbf{T} satisfies detailed balance relative to $\boldsymbol{\pi}$ if and only if \mathbf{T} is self-adjoint under the inner product $\langle \cdot, \cdot \rangle_{\boldsymbol{\pi}}$.*

Proof of Proposition 3: Assume that detailed balance is satisfied. Take $\mathbf{v} = \boldsymbol{\delta}_i$ and $\mathbf{w} = \boldsymbol{\delta}_j$ for some $i, j \in N$, i.e. take two standard basis vectors. Then

$$\langle \mathbf{T}\mathbf{v}, \mathbf{w} \rangle_{\boldsymbol{\pi}} = T_{ji}\pi_j$$

and

$$\langle \mathbf{v}, \mathbf{T}\mathbf{w} \rangle_{\boldsymbol{\pi}} = T_{ij}\pi_i.$$

These two quantities are equal by detailed balance and this equality extends to arbitrary \mathbf{v} and \mathbf{w} because the inner product is a bilinear form.

For the converse direction, the equality $\langle \mathbf{T}\boldsymbol{\delta}_i, \boldsymbol{\delta}_j \rangle_{\boldsymbol{\pi}} = \langle \boldsymbol{\delta}_i, \mathbf{T}\boldsymbol{\delta}_j \rangle_{\boldsymbol{\pi}}$ for standard basis vectors is guaranteed as a consequence of \mathbf{T} being self-adjoint, and this immediately gives detailed balance by the simple calculation above. ■

The next claim is that $\mathbf{T}(\mathbf{A})$ satisfies detailed balance relative to $\mathbf{s}(\mathbf{A})$, which is defined by

$$s_i(\mathbf{A}) = \frac{d_i(\mathbf{A})}{\sum_i d_i(\mathbf{A})}.$$

This is immediate to check from the definitions. Thus, $\mathbf{s}(\mathbf{A})$ is the stationary distribution of $\mathbf{T}(\mathbf{A})$ and, moreover, $\mathbf{T}(\mathbf{A})$ is self-adjoint relative to $\langle \cdot, \cdot \rangle_{\mathbf{s}(\mathbf{A})}$. As a result, the eigenvalues of $\mathbf{T}(\mathbf{A})$ are all real. Let $1 = \lambda_1(\mathbf{T}(\mathbf{A})), \dots, \lambda_n(\mathbf{T}(\mathbf{A}))$ denote these eigenvalues ordered from greatest to least by magnitude, and

$$1 = \beta_1(\mathbf{T}(\mathbf{A})) > \beta_2(\mathbf{T}(\mathbf{A})) \geq \beta_3(\mathbf{T}(\mathbf{A})) \geq \dots \geq \beta_n(\mathbf{T}(\mathbf{A})) > -1$$

denote these same eigenvalues ordered from greatest to least as real numbers. Obviously

$$|\lambda_2(\mathbf{T}(\mathbf{A}))| = \max\{|\beta_2(\mathbf{T}(\mathbf{A}))|, |\beta_n(\mathbf{T}(\mathbf{A}))|\}.$$

Now, since $\mathbf{T}(\mathbf{A})$ is self-adjoint, we can use the powerful Courant-Fischer variational characterization of the eigenvalues of a Hermitian matrix¹⁶. Let \mathbf{e} denote the unit column vector of ones.

PROPOSITION 4.

$$\beta_n(\mathbf{T}(\mathbf{A})) = \inf_{\mathbf{0} \neq \mathbf{v} \in \mathbb{R}^n} \left\{ \frac{\langle \mathbf{v}, \mathbf{T}(\mathbf{A})\mathbf{v} \rangle}{\langle \mathbf{v}, \mathbf{v} \rangle} \right\}. \quad (7)$$

$$\beta_2(\mathbf{T}(\mathbf{A})) = \sup_{\substack{\mathbf{0} \neq \mathbf{v} \in \mathbb{R}^n \text{ s.t.} \\ \langle \mathbf{v}, \mathbf{e} \rangle = 0}} \left\{ \frac{\langle \mathbf{v}, \mathbf{T}(\mathbf{A})\mathbf{v} \rangle}{\langle \mathbf{v}, \mathbf{v} \rangle} \right\}, \quad (8)$$

where the inner product everywhere is $\langle \cdot, \cdot \rangle_{\mathbf{s}(\mathbf{A})}$.

This says that the smallest eigenvalue (under the real number ordering) minimizes the normalized quadratic form in braces, where \mathbf{v} ranges over all nonzero vectors. Moreover, the second largest eigenvalue (under the real number ordering) maximizes the same quantity but where \mathbf{v} ranges over all nonzero vectors orthogonal to the right eigenvector corresponding to the largest eigenvalue.

7.2 Proofs and Additional Material for the Main Results

7.2.1 Relating Consensus Times to Second Eigenvalues

Proof of Lemma 3: In the proof, we fix \mathbf{A} and drop it as an argument; we also drop the argument on the eigenvalues, being fixed throughout.

We first show that

$$\text{CT}(\varepsilon) \leq \left\lceil \frac{\log(1/\varepsilon)}{2 \log(1/|\lambda_2|)} \right\rceil.$$

Take any $\mathbf{b} \in [0, 1]^n$. Let \mathbf{U}_i be the projection onto the eigenspace of \mathbf{T} corresponding to λ_i . Note that under $\langle \cdot, \cdot \rangle_{\mathbf{s}}$, these eigenspaces are orthogonal. Define $\mathbf{U} = \sum_{i=2}^n \mathbf{U}_i$. This is the projection off the eigenspace corresponding to $\lambda = 1$. Then:

$$\begin{aligned} \|(\mathbf{T}^t - \mathbf{T}^\infty)\mathbf{b}\|_{\mathbf{s}}^2 &= \left\| \sum_{i=2}^n \lambda_i^t \mathbf{U}_i \mathbf{b} \right\|_{\mathbf{s}}^2 && \text{spectral theorem applied to the stochastic matrix} \\ &= \sum_{i=2}^n |\lambda_i|^{2t} \|\mathbf{U}_i \mathbf{b}\|_{\mathbf{s}}^2 && \text{orthogonality of the spectral projections} \\ &\leq |\lambda_2|^{2t} \sum_{i=2}^n \|\mathbf{U}_i \mathbf{b}\|_{\mathbf{s}}^2 \end{aligned}$$

¹⁶Horn and Johnson (1985, p. 176–178).

$$\begin{aligned}
&= |\lambda_2|^{2t} \left\| \sum_{i=2}^n \mathbf{U}_i \mathbf{b} \right\|_{\mathbf{s}}^2 && \text{orthogonality of the spectral projections} \\
&= |\lambda_2|^{2t} \|\mathbf{U}\mathbf{b}\|_{\mathbf{s}}^2 && \text{definition of } \mathbf{U} \\
&\leq |\lambda_2|^{2t} \|\mathbf{b}\|_{\mathbf{s}}^2 && \text{projections are contractions} \\
&\leq |\lambda_2|^{2t} \sum_{i=1}^n s_i && \mathbf{b} \in [0, 1]^n \text{ and definition of } \langle \cdot, \cdot \rangle_{\mathbf{s}} \\
&= |\lambda_2|^{2t} && \text{entries of } \mathbf{s} \text{ sum to } 1.
\end{aligned}$$

Thus, if

$$t \geq \frac{\log(1/\varepsilon)}{2 \log(1/|\lambda_2|)}$$

then

$$\|(\mathbf{T}^t - \mathbf{T}^\infty)\mathbf{b}\|_{\mathbf{s}}^2 \leq \varepsilon,$$

from which the bound follows upon observing that $\text{CT}(\varepsilon)$ must be an integer. This also shows that when the second eigenvalue is identically 0, then consensus time must be 1.

Now we show that

$$\left\lfloor \frac{\log(\underline{s}/4\varepsilon)}{2 \log(1/|\lambda_2|)} \right\rfloor \leq \text{CT}(\varepsilon).$$

Let \mathbf{w} be an eigenvector of \mathbf{T} corresponding to λ_2 , scaled so that $\|\mathbf{w}\|_{\mathbf{s}}^2 = \underline{s}/4$. Then the maximum entry of \mathbf{w} is at most $1/2$ and the minimum entry is at least $-1/2$. Consequently, if we define $\mathbf{b} = \mathbf{w} + \mathbf{e}/2$, then $\mathbf{b} \in [0, 1]^n$. Now, using the fact that \mathbf{e} is a right eigenvector corresponding to $\lambda_1 = 1$ and spectral projections are orthogonal, it follows that:

$$\begin{aligned}
\|(\mathbf{T}^t - \mathbf{T}^\infty)\mathbf{b}\|_{\mathbf{s}}^2 &= |\lambda_2|^{2t} \|\mathbf{U}_2 \mathbf{w}\|_{\mathbf{s}}^2 \\
&= |\lambda_2|^{2t} \|\mathbf{w}\|_{\mathbf{s}}^2 \\
&= \frac{\underline{s}}{4} |\lambda_2|^{2t}.
\end{aligned}$$

Therefore, if

$$t \leq \frac{\log(\underline{s}/4\varepsilon)}{2 \log(1/|\lambda_2|)}$$

then

$$\|(\mathbf{T}^t - \mathbf{T}^\infty)\mathbf{b}\|_{\mathbf{s}}^2 \geq \varepsilon,$$

from which the remaining bound follows upon observing that $\text{CT}(\varepsilon)$ must be an integer. \blacksquare

7.2.2 The Representative Agent Theorem and a Consequence

Theorem 3 and Proposition 6 require related machinery, which we will develop and apply in this section. First, we introduce some notation.

We drop the arguments on the random variable \mathbf{A} . Let $\mathbf{D}(\mathbf{A})$ denote the diagonal matrix whose (i, i) entry is $d_i(\mathbf{A})$. Let \mathbf{R} be the n -by- n matrix given by $R_{ij} = P_{kl}$ if $i \in N_k, j \in N_\ell$. Then the expected degree of node i is $w_i := \sum_j R_{ij}$.

Let $V = \sum_i w_i$ be the sum of expected degrees and $v = \sum_i d_i(\mathbf{A})$ the sum of realized degrees.

For any matrix \mathbf{T} , let $\|\mathbf{T}\| = \sup_{\|\mathbf{v}\|=1} \langle \mathbf{v}, \mathbf{T}\mathbf{v} \rangle$, where the inner product is the standard Euclidean dot product. Let

$$\mathbf{J} = \mathbf{D}(\mathbf{A})^{-1/2} \mathbf{A} \mathbf{D}(\mathbf{A})^{-1/2} - v^{-1} \mathbf{D}(\mathbf{A})^{1/2} \mathbf{E} \mathbf{D}(\mathbf{A})^{1/2}$$

and

$$\mathbf{K} = \mathbf{D}(\mathbf{R})^{-1/2} \mathbf{R} \mathbf{D}(\mathbf{R})^{-1/2} - V^{-1} \mathbf{D}(\mathbf{R})^{1/2} \mathbf{E} \mathbf{D}(\mathbf{R})^{1/2}.$$

Now we note a fact from basic linear algebra.

FACT 1. $\mathbf{D}(\mathbf{A})^{-1/2} \mathbf{A} \mathbf{D}(\mathbf{A})^{-1/2}$ and $\mathbf{T}(\mathbf{A}) = \mathbf{D}(\mathbf{A})^{-1} \mathbf{A}$ are similar matrices, so that they have the same eigenvalues, and that $v^{-1} \mathbf{D}(\mathbf{A})^{1/2} \mathbf{E} \mathbf{D}(\mathbf{A})^{1/2}$ is the summand of the spectral decomposition of $\mathbf{D}(\mathbf{A})^{-1/2} \mathbf{A} \mathbf{D}(\mathbf{A})^{-1/2}$ corresponding to the eigenvalue 1. The same reasoning applies when we replace \mathbf{A} by \mathbf{R} and v by V .

We now state the proof of Theorem 3, or rather a reduction to a proposition which will also be useful for proving Proposition 6.

Proof of Theorem 3: It is clear that $\mathbf{D}(\mathbf{R})^{-1} \mathbf{R}$ has the same eigenvalues as \mathbf{Q} , so to prove the claim it suffices to prove that the former matrix has second eigenvalue close enough to that of $\mathbf{D}(\mathbf{A})^{-1/2} \mathbf{A} \mathbf{D}(\mathbf{A})^{-1/2}$.

By Fact 1, we know that $\|\mathbf{J}\|$ is the second largest eigenvalue in magnitude of the matrix $\mathbf{D}(\mathbf{A})^{-1/2} \mathbf{A} \mathbf{D}(\mathbf{A})^{-1/2}$, and $\|\mathbf{K}\|$ is the second largest eigenvalue in magnitude of the matrix $\mathbf{D}(\mathbf{R})^{-1/2} \mathbf{R} \mathbf{D}(\mathbf{R})^{-1/2}$. Thus, by the triangle inequality, if we can show that with probability at least $1 - \delta$ we have $\|\mathbf{J} - \mathbf{K}\| < \delta$, then the proof is done. This is the content of Proposition 5 below. ■

Now we state a lemma from the proof of Theorem 3.6 of Chung Lu and Vu (2004), which is a consequence of a Chernoff-type concentration inequality and is quite useful throughout this section.

LEMMA 7. Fix any $\delta > 0$. If $w_{\min}/\log n$ is high enough, the following statement holds with probability at least $1 - \delta$ for all i simultaneously: $|d_i - w_i| < \delta w_i$.

PROPOSITION 5. If $w_{\min}/\log^2 n$ is high enough, then with probability at least $1 - \delta$ we have $\|\mathbf{J} - \mathbf{K}\| < \delta$.

Proof of Proposition 5: Write

$$\mathbf{J} - \mathbf{K} = \mathbf{B} + \mathbf{C} + \mathbf{L} + \mathbf{M} \quad \text{where} \quad \begin{aligned} B_{ij} &= \frac{A_{ij}}{\sqrt{d_i d_j}} \left(1 - \frac{\sqrt{d_i d_j}}{\sqrt{w_i w_j}} \right) \\ C_{ij} &= \frac{A_{ij} - R_{ij}}{\sqrt{w_i w_j}} \\ L_{ij} &= \frac{\sqrt{w_i w_j}}{V} \left(1 - \frac{\sqrt{d_i d_j}}{\sqrt{w_i w_j}} \right) \\ M_{ij} &= (V^{-1} - v^{-1}) \sqrt{d_i d_j}. \end{aligned}$$

By the triangle inequality,

$$\|\mathbf{J} - \mathbf{K}\| \leq \|\mathbf{B}\| + \|\mathbf{C}\| + \|\mathbf{L}\| + \|\mathbf{M}\|,$$

so it suffices to bound the pieces individually.

Now we list two lemmas, useful only in this proof, from Chung Lu and Vu (2004). The proof of the first requires only minor modification in our setting.

LEMMA 8. Fix any $\delta > 0$. Then if $w_{\min}/\log^2 n$ is high enough, with probability at least $1 - \delta$:

$$\|\mathbf{C}\| \leq \frac{2}{\sqrt{\bar{w}}} + \frac{\log n}{\sqrt{w_{\min}}}.$$

Proof of Lemma 8: The only step of the proof of this last lemma that does not work exactly as in the proofs of Theorems 3.2 and 3.6 of Chung Lu and Vu (2004) is their equation (3.2). This step asserts (in our notation) that for $m \geq 2$, we have

$$\mathbb{E}[C_{ij}^m] \leq \frac{(1 - R_{ij})R_{ij} + (-R_{ij})^m(1 - R_{ij})}{(w_i w_j)^{m/2}} \leq \frac{R_{ij}}{(w_i w_j)^{m/2}} \leq \frac{w_i w_j / V}{(w_i w_j)^{m/2}} \leq \frac{1/V}{(w_{\min})^{m-2}}.$$

The step which is slightly different is the penultimate inequality. We must show that $R^{ij} \leq w_i w_j / V$. But note that $w_j / V \leq 1$ by definition and $w_i = \sum_k R_{ik} \geq R_{ij}$. \blacksquare

It follows that we have $\|\mathbf{C}\| < \delta/4$ with probability at least $1 - \delta/4$.

LEMMA 9. *Fix any $\delta > 0$. If $w_{\min}/\log^2 n$ is high enough, the following statement holds with probability at least $1 - \delta$:*

$$\|\mathbf{M}\| \leq \frac{1}{\sqrt{w}}.$$

It follows that we have $\|\mathbf{M}\| < \delta/4$ with probability at least $1 - \delta/4$.

To bound $\|\mathbf{B}\|$ and $\|\mathbf{L}\|$, we will use Lemma 7 and two simple facts about the matrix norm. Let $\text{abs}(\mathbf{X})$ denote the matrix whose (i, j) entry is $|X_{ij}|$.

LEMMA 10. 1. *For any matrix \mathbf{X} , $\|\mathbf{X}\| \leq \|\text{abs}(\mathbf{X})\|$.*

2. *Suppose there are two nonnegative matrices, \mathbf{X} and \mathbf{Y} and a constant $c > 0$ such that for each i, j , we have $Y_{ij} < cX_{ij}$. Then $\|\mathbf{Y}\| \leq c\|\mathbf{X}\|$.*

Proof of Lemma 10: For (1), note that for all $\|\mathbf{v}\| = 1$, we have

$$\begin{aligned} \langle \mathbf{v}, \mathbf{X}\mathbf{v} \rangle &= \sum_{i,j} v_i v_j X_{ij} \\ &\leq \sum_{i,j} |v_i v_j X_{ij}| \\ &\leq \sum_{i,j} |v_i| |v_j| |X_{ij}| \\ &\leq \|\text{abs}(\mathbf{X})\|, \end{aligned}$$

the last inequality being true because $\|\text{abs}(\mathbf{v})\| = 1$. This proves the claim by definition of the matrix norm.

For (2), note that for all $\|\mathbf{v}\| = 1$, we have

$$\begin{aligned} \langle \mathbf{v}, \mathbf{Y}\mathbf{v} \rangle &= \sum_{i,j} v_i v_j Y_{ij} \\ &< \sum_{i,j} |v_i| |v_j| c X_{ij} \\ &\leq c \sum_{i,j} v_i v_j X_{ij} \\ &\leq c\|\mathbf{X}\|, \end{aligned}$$

where again we have made use of the fact that $\|\text{abs}(\mathbf{v})\| = 1$. ■

To show that, with probability at least $1 - \delta$, we have $\|\mathbf{B}\| < \delta/4$, define $\hat{\mathbf{B}} = \text{abs}(\mathbf{B})$; by Lemma 10(1) it suffices to show $\|\hat{\mathbf{B}}\| < \delta/4$. Note

$$\hat{B}_{ij} = \frac{A_{ij}}{\sqrt{d_i d_j}} \left| 1 - \frac{\sqrt{d_i d_j}}{\sqrt{w_i w_j}} \right|.$$

By Lemma 7 we have with probability at least $1 - \delta/4$ that

$$\left| 1 - \frac{\sqrt{d_i d_j}}{\sqrt{w_i w_j}} \right| < \delta/4$$

and so, noting that

$$\|\mathbf{D}(\mathbf{A})^{-1/2} \mathbf{A} \mathbf{D}(\mathbf{A})^{-1/2}\| = 1$$

and using Lemma 10(2), the claim is proved.

Precisely the same argument works to show that with probability at least $1 - \delta/4$, we have $\|\mathbf{L}\| < \delta/4$, with $V^{-1} \mathbf{D}(\mathbf{R})^{1/2} \mathbf{E} \mathbf{D}(\mathbf{R})^{1/2}$, which also has norm 1, playing the role of $\mathbf{D}(\mathbf{A})^{-1/2} \mathbf{A} \mathbf{D}(\mathbf{A})^{-1/2}$.

Combining all the bounds shows that, with probability at least $1 - \delta$ we have $\|\mathbf{J} - \mathbf{K}\| < \delta$, as desired.

This completes the proof of the proposition. ■

We will now use the results established so far in this section to prove a proposition that tightens the lower bound in Lemma 3, so that second eigenvalues become an even better proxy for consensus time.

PROPOSITION 6. *Suppose (\mathbf{P}, \mathbf{n}) are such that, for all n ,*

1. *There exist $\underline{\lambda}$ and $\bar{\lambda}$ so that $0 < \underline{\lambda} \leq \lambda_2(\mathbf{Q}(\mathbf{P}, \mathbf{n})) \leq \bar{\lambda} < 1$.*
2. *$\min_k n_k/n \geq \alpha > 0$.*
3. *$d_{\min}(\mathbf{P})/d_{\max}(\mathbf{P}) \geq \beta > 0$.*

Write $\mathbf{T} = \mathbf{T}(\mathbf{A})$. Then, for any $\delta > 0$, for high enough n , with probability at least $1 - \delta$

$$\left\lfloor \frac{\log(1/8\varepsilon) - \log(1/\alpha\beta)}{2 \log(1/|\lambda_2(\mathbf{T})|)} \right\rfloor - 1 \leq \text{CT}(\varepsilon; \mathbf{A}).$$

Combining this with Lemma 3, we can conclude that for any δ with probability at least $1 - \delta$

$$\left\lfloor \frac{\log(1/8\varepsilon) - \log(1/\alpha\beta)}{2 \log(1/|\lambda_2(\mathbf{T})|)} \right\rfloor - 1 \leq \text{CT}(\varepsilon; \mathbf{A}) \leq \left\lceil \frac{\log(1/\varepsilon)}{2 \log(1/|\lambda_2(\mathbf{T})|)} \right\rceil.$$

Thus, as we let ε get small, we find that $\text{CT}(\varepsilon; \mathbf{A})$ is proportional to $\frac{\log(1/\varepsilon)}{\log(1/|\lambda_2(\mathbf{T})|)}$.

The assumptions of this proposition could be weakened, but in their current form they are simple to state and interpret. The first one says that the second eigenvalue of $\mathbf{Q}(\mathbf{P}, \mathbf{n})$ should have a magnitude that stays away from 0 and 1, and amounts to requiring that consensus time is not going to 0 or ∞ . The other conditions impose some balance on the system. The second one says that no group should be getting negligibly small relative to society. The third one says that maximum and minimum degrees should not get too far apart proportionally. Some types are allowed to be much more popular than others, but not infinitely so. If these conditions are met, then Lemma 3 can be strengthened so that the lower bound is tighter, and still easy to compute.

Techniques similar to the ones used in the proof below can be applied to the study of mixing times in order to tighten the upper bound in Lemma 4 in the multi-type random graph setting.

Proof of Proposition 6: We will reuse the same variable names used inside the proof of Proposition 5, but the variables defined for the whole subsection will be unchanged.

Write $\mathbf{C} = \mathbf{D}(\mathbf{R})^{-1}\mathbf{R}$ and $\mathbf{T} = \mathbf{T}(\mathbf{A})$. That is, \mathbf{C} is the version of \mathbf{T} in the “expectations” world. Also, let

$$z = \left\lfloor \frac{\log(1/8\varepsilon) - \log(1/\alpha\beta)}{2 \log(1/|\lambda_2(\mathbf{T})|)} \right\rfloor.$$

There are three steps to the proof. In Step 1, we show that for $\mathbf{C}^t \mathbf{b}$ to converge within 2ε of its limit takes at least $z - 1$ steps for some \mathbf{b} . In Step 2, we use Proposition 5 to show that for any $\eta > 0$, for high enough n , with probability at least $1 - \eta$, we have $\|\mathbf{T} - \mathbf{C}\| < \eta$. In Step 3, we show that, if η is chosen small enough, then $\mathbf{C}^t \mathbf{b}$ and $\mathbf{T}^t \mathbf{b}$ are at most ε apart for at least $z - 1$ steps under the inner product $\langle \cdot, \cdot \rangle_{\mathbf{s}(\mathbf{A})}$. This shows the requisite result.

Step 1. Let \mathbf{v} be a right eigenvector of \mathbf{C} corresponding to eigenvalue $\hat{\lambda}_2 := \lambda_2(\mathbf{Q})$ (this is also the second eigenvalue in magnitude of \mathbf{C} by Fact 1). If we multiply \mathbf{v} by a constant scalar, we may assume that the entry with largest magnitude is $1/2$. By Assumption 1, $\hat{\lambda}_2$ is nonzero. Given this and the fact that \mathbf{C} is constant on a given type, it follows that \mathbf{v} is constant on a given type. Thus, by Assumption 2, there are at least αn entries in \mathbf{v} equal

to $1/2$. And from this it follows, by the definition of $\mathbf{s}(\mathbf{C})$ and Assumption 3, that

$$\langle \mathbf{v}, \mathbf{v} \rangle_{\mathbf{s}(\mathbf{C})} \geq n\alpha \cdot \left(\frac{1}{2}\right)^2 \cdot \frac{d_{\min}(\mathbf{P})}{nd_{\max}(\mathbf{P})} \geq \frac{\alpha\beta}{4}.$$

Setting $b_i = v_i + 1/2$, we see as at the end of the proof of Lemma 3 that

$$\|\mathbf{C}^t \mathbf{b} - \mathbf{C}^\infty \mathbf{b}\|_{\mathbf{s}(\mathbf{C})} \geq \frac{\alpha\beta}{4} |\lambda_2(\mathbf{C})|^{2t},$$

which yields the lower bound on convergence time we want with \mathbf{C} instead of \mathbf{T} . But in view of Assumption 1 and Theorem 3, for high enough n we can replace \mathbf{C} by \mathbf{T} and lose at most an additive factor of 1 in the bound.

Step 2. Recall $\mathbf{C} = \mathbf{D}(\mathbf{R})^{-1} \mathbf{R}$ and $\mathbf{T} = \mathbf{T}(\mathbf{A})$. Also put

$$\mathbf{L} = \mathbf{D}(\mathbf{A})^{-1/2} \mathbf{J} \mathbf{D}(\mathbf{A})^{1/2}$$

and

$$\mathbf{M} = \mathbf{D}(\mathbf{R})^{-1/2} \mathbf{K} \mathbf{D}(\mathbf{R})^{1/2}.$$

By Fact 1, we have

$$\mathbf{T} - \mathbf{C} = v^{-1} \mathbf{E} \mathbf{D}(\mathbf{A}) - V^{-1} \mathbf{E} \mathbf{D}(\mathbf{R}) + \mathbf{L} - \mathbf{M}.$$

So by the triangle inequality, it suffices to bound $\|v^{-1} \mathbf{E} \mathbf{D}(\mathbf{A}) - V^{-1} \mathbf{E} \mathbf{D}(\mathbf{R})\|$ and $\|\mathbf{L} - \mathbf{M}\|$. By Lemma 7, if $w_{\min}/\log^2 n$ is high enough, the following event occurs with probability at least $1 - \gamma$ for all i simultaneously: $|d_i - w_i| < \gamma w_i$. Call this event E_1 . Given the assumptions, high enough n ensures the condition of the lemma is met. Thus, on E_1 ,

$$\|v^{-1} \mathbf{E} \mathbf{D}(\mathbf{A}) - V^{-1} \mathbf{E} \mathbf{D}(\mathbf{R})\| < \gamma,$$

and so it suffices to take care of the other term.

By Proposition 5, we know that if $w_{\min}/\log^2 n$ is high enough, then on an event E_2 of probability at least $1 - \gamma$ we have $\|\mathbf{J} - \mathbf{K}\| < \gamma$. As above, for high enough n the condition is met. Now let

$$\mathbf{F} = \mathbf{D}(\mathbf{A}) - \mathbf{D}(\mathbf{R}),$$

$$\mathbf{G} = (\mathbf{D}(\mathbf{R}) + \mathbf{F})^{1/2} - \mathbf{D}(\mathbf{R})^{1/2},$$

and

$$\mathbf{H} = (\mathbf{D}(\mathbf{R}) + \mathbf{F})^{-1/2} - \mathbf{D}(\mathbf{R})^{-1/2}.$$

Observe

$$\begin{aligned} \|\mathbf{L} - \mathbf{M}\| &= \|(\mathbf{D}(\mathbf{R}) + \mathbf{F})^{-1/2}\mathbf{J}(\mathbf{D}(\mathbf{R}) + \mathbf{F})^{1/2} - \mathbf{D}(\mathbf{R})^{-1/2}\mathbf{K}\mathbf{D}(\mathbf{R})^{1/2}\| \\ &= \|(\mathbf{D}(\mathbf{R})^{-1/2} + \mathbf{H})\mathbf{J}(\mathbf{D}(\mathbf{R})^{1/2} + \mathbf{G}) - \mathbf{D}(\mathbf{R})^{-1/2}\mathbf{K}\mathbf{D}(\mathbf{R})^{1/2}\| \\ &= \|\mathbf{D}(\mathbf{R})^{-1/2}(\mathbf{J} - \mathbf{K})\mathbf{D}(\mathbf{R})^{1/2} + \mathbf{D}(\mathbf{R})^{-1/2}\mathbf{J}\mathbf{G} + \mathbf{H}\mathbf{J}\mathbf{D}(\mathbf{R})^{1/2}\| \\ &\leq \|\mathbf{D}(\mathbf{R})^{-1/2}(\mathbf{J} - \mathbf{K})\mathbf{D}(\mathbf{R})^{1/2}\| + \|\mathbf{D}(\mathbf{R})^{-1/2}\mathbf{J}\mathbf{G}\| + \|\mathbf{H}\mathbf{J}\mathbf{D}(\mathbf{R})^{1/2}\| + \|\mathbf{H}\mathbf{J}\mathbf{G}\| \end{aligned}$$

Using Lemma 7 and standard series approximation arguments, for high enough n we can ensure $\|\mathbf{G}\| \leq \gamma\|\mathbf{D}(\mathbf{R})^{1/2}\|$ and $\|\mathbf{H}\| \leq \gamma\|\mathbf{D}(\mathbf{R})^{-1/2}\|$ on an event E_3 of probability at least $1 - \gamma$. Using the fact that $\|\mathbf{J}\| \leq 1$, the Cauchy-Schwartz inequality yields that each of the middle two terms above is bounded by γ . For the last term, note that

$$\|\mathbf{H}\mathbf{J}\mathbf{G}\| \leq \gamma^2\|\mathbf{D}(\mathbf{R})^{1/2}\| \cdot \|\mathbf{D}(\mathbf{R})^{-1/2}\| = \frac{d_{\max}(\mathbf{P})}{d_{\min}(\mathbf{P})} \leq \frac{\gamma^2}{\beta}.$$

So it suffices to take care of the first term. This is accomplished by noticing that, on $E_1 \cap E_2$,

$$\begin{aligned} \|\mathbf{D}(\mathbf{R})^{-1/2}(\mathbf{J} - \mathbf{K})\mathbf{D}(\mathbf{R})^{1/2}\| &\leq \frac{d_{\max}(\mathbf{A})^{1/2}}{d_{\min}(\mathbf{A})^{1/2}}\|\mathbf{J} - \mathbf{K}\| \\ &\leq (1 + \gamma)\frac{d_{\max}(\mathbf{P})^{1/2}}{d_{\min}(\mathbf{P})^{1/2}}\|\mathbf{J} - \mathbf{K}\| && \text{definition of } E_1 \\ &\leq \frac{1 + \gamma}{\beta}\|\mathbf{J} - \mathbf{K}\| && \text{Assumption 3} \\ &\leq \frac{(1 + \gamma)\gamma}{\beta} && \text{definition of } E_2. \end{aligned}$$

Together, these facts show that for high enough n , on $E_1 \cap E_2 \cap E_3$, which occurs with probability at least $1 - 3\gamma$, we have

$$\|\mathbf{T} - \mathbf{C}\| \leq \gamma + \frac{(1 + \gamma)\gamma}{\beta} + 2\gamma + \frac{\gamma^2}{\beta}.$$

By choosing γ so that the right hand side is less than η and $3\gamma < \eta$ (to take care of the probability), the step is complete.

Step 3. Write $\mathbf{T} = \mathbf{C} + \mathbf{Y}$, where $\|\mathbf{Y}\| \leq \eta$. Note that

$$(\mathbf{T} + \mathbf{Y})^t = \mathbf{T}^t + \sum_{q=0}^{t-1} \mathbf{X}_q,$$

where \mathbf{X}_q is a product of q copies of \mathbf{Y} and $t - q$ copies of \mathbf{T} in some order. By the fact that $\|\mathbf{T}\| = 1$ and $\|\mathbf{Y}\| \leq \eta$, we have $\|\mathbf{X}_q\| \leq \eta^q$ for each $q \geq 1$. Then, by the triangle inequality,

$$\left\| \sum_{q=0}^{t-1} \mathbf{X}_q \right\| \leq \sum_{q=0}^{t-1} \eta^q \leq \frac{1}{1 - \eta}.$$

Thus,

$$\mathbf{Y}_t := \|\mathbf{C}^t - \mathbf{T}^t\| \leq \frac{1}{1 - \eta}.$$

Take \mathbf{b} and \mathbf{v} to be the vectors constructed in Step 1. Note that for $t \leq z - 1$ we have, for high enough n ,

$$\begin{aligned} \langle \mathbf{T}^t \mathbf{b} - \mathbf{T}^\infty \mathbf{b}, \mathbf{T}^t \mathbf{b} - \mathbf{T}^\infty \mathbf{b} \rangle_{\mathbf{s}(\mathbf{A})} &= \langle \mathbf{T}^t \mathbf{v}, \mathbf{T}^t \mathbf{v} \rangle_{\mathbf{s}(\mathbf{A})} \\ &= \langle (\mathbf{C}^t + \mathbf{Y}_t) \mathbf{v}, (\mathbf{C}^t + \mathbf{Y}_t) \mathbf{v} \rangle_{\mathbf{s}(\mathbf{A})} \\ &\geq \langle \mathbf{C}^t \mathbf{v}, \mathbf{C}^t \mathbf{v} \rangle_{\mathbf{s}(\mathbf{A})} + 2 \langle \mathbf{Y}_t \mathbf{v}, \mathbf{C}^t \mathbf{v} \rangle_{\mathbf{s}(\mathbf{A})} \\ &\geq (1 - \eta) \langle \mathbf{C}^t \mathbf{v}, \mathbf{C}^t \mathbf{v} \rangle_{\mathbf{s}(\mathbf{C})} \\ &\quad + 2 \langle \mathbf{Y}_t \mathbf{v}, \mathbf{C}^t \mathbf{v} \rangle_{\mathbf{s}(\mathbf{A})} && \text{Lemma 7} \\ &\leq 2(1 - \eta)\varepsilon + 2 \langle \mathbf{Y}_t \mathbf{v}, \mathbf{C}^t \mathbf{v} \rangle_{\mathbf{s}(\mathbf{A})} && \text{definition of } z \\ &\leq 2(1 - \eta)\varepsilon - 2 \|\mathbf{Y}_t \mathbf{v}\|_{\mathbf{s}(\mathbf{A})} \cdot \|\mathbf{C}^t \mathbf{v}\|_{\mathbf{s}(\mathbf{A})} && \text{Cauchy-Schwartz} \\ &\leq 2(1 - \eta)\varepsilon - 2 \|\mathbf{Y}_t \mathbf{v}\|_{\mathbf{s}(\mathbf{A})} && \text{see below} \\ &\leq 2(1 - \eta)\varepsilon - 2 \|\mathbf{Y}_t \mathbf{v}\| && \text{def'n of } \|\cdot\|_{\mathbf{s}(\mathbf{A})} \\ &\leq 2(1 - \eta)\varepsilon - 2\eta. \end{aligned}$$

The step whose explanation is missing is straightforward; no entries in \mathbf{v} have magnitude exceeding $1/2$ and multiplication by the stochastic matrix \mathbf{C} preserves this property. Since $\mathbf{s}(\mathbf{C})$ is a probability distribution, $\|\mathbf{C}^t \mathbf{v}\|_{\mathbf{s}(\mathbf{C})} \leq 1$ holds by definition of the inner product. If η is chosen so that $2(1 - \eta)\varepsilon - 2\eta > \varepsilon$, then the proof is complete. \blacksquare

7.2.3 Results on DWH and EDWH

Proof of Proposition 1: We will construct a \mathbf{v} satisfying $\langle \mathbf{v}, \mathbf{e} \rangle_{\mathbf{s}} = 0$ so that the absolute value of the quantity $\langle \mathbf{v}, \mathbf{T}\mathbf{v} \rangle_{\mathbf{s}} / \langle \mathbf{v}, \mathbf{v} \rangle_{\mathbf{s}}$ is equal to $|\text{DWH}(M)|$. Since $|\lambda_2| = \max\{|\beta_2|, |\beta_n|\}$, this suffices by Proposition 4.

Define

$$v_i = \begin{cases} \frac{1}{rd_i} & \text{if } i \in M \\ -\frac{1}{(n-r)d_i} & \text{if } i \notin M. \end{cases}$$

Let $D = \sum_i d_i$ and note

$$\begin{aligned} \langle \mathbf{v}, \mathbf{v} \rangle_{\mathbf{s}} &= \sum_i v_i^2 \cdot s_i \\ &= \sum_{i \in M} \frac{1}{(rd_i)^2} \cdot \frac{d_i}{D} + \sum_{i \in M^c} \frac{1}{((n-r)d_i)^2} \cdot \frac{d_i}{D} \\ &= \frac{1}{D} \left[\frac{1}{r^2} \sum_{i \in M} \frac{1}{d_i} + \frac{1}{(n-r)^2} \sum_{i \in M^c} \frac{1}{d_i} \right]. \end{aligned} \quad (9)$$

Also,

$$\begin{aligned} \langle \mathbf{v}, \mathbf{T}\mathbf{v} \rangle_{\mathbf{s}} &= \sum_i v_i \left(\sum_j T_{ij} v_j \right) s_i \\ &= \frac{1}{D} \sum_i \sum_j v_i T_{ij} v_j d_i \\ &= \frac{1}{D} \sum_{i,j:T_{ij}>0} v_i \frac{1}{d_i} v_j d_i \\ &= \frac{1}{D} \left[\frac{1}{r^2} \sum_{i,j \in M} T_{ij} T_{ji} + \frac{1}{(n-r)^2} \sum_{i,j \in M^c} T_{ij} T_{ji} - \frac{2}{r(n-r)} \sum_{i \in M, j \in M^c} T_{ij} T_{ji} \right]. \end{aligned}$$

Dividing $\langle \mathbf{v}, \mathbf{T}\mathbf{v} \rangle_{\mathbf{s}}$ by $\langle \mathbf{v}, \mathbf{v} \rangle_{\mathbf{s}}$, canceling D , and using the definition of W yields the result. \blacksquare

We prove Lemma 5 before Theorem 2.

Proof of Lemma 5: We show that $\text{EDWH}(m, p_s, p_d) = (p_s - p_d)/mp = (p_s - p_d)/(p_s + (m-1)p_d)$, which is easily checked to be equal to $(H-1)/(m-1)$; this, in turn, converges to $h = H/m$ as m grows. Consider M consisting of k islands and M^c consisting of $m-k$

islands of nodes. Then from the definition of $\text{EDWH}(M; m, p_s, p_d)$ it follows that

$$\text{EDWH}(M; m, p_s, p_d) = \frac{EW_{M,M} + EW_{M^c, M^c} - 2EW_{M, M^c}}{\frac{1}{|M|^2} \sum_{i \in M} \frac{1}{d} + \frac{1}{|M^c|^2} \sum_{i \in M^c} \frac{1}{d}},$$

which can be written as

$$\frac{\frac{kp_s + k(k-1)p_d}{k^2 d^2} + \frac{(m-k)p_s + (m-k)(m-k-1)p_d}{(m-k)^2 d^2} - 2\frac{p_d}{d^2}}{\frac{m^2}{k^2 n^2} \frac{kn}{dm} + \frac{m^2}{(m-k)^2 n^2} \frac{(m-k)n}{dm}}.$$

This becomes

$$\frac{\frac{p_s + (k-1)p_d}{k} + \frac{p_s + (m-k-1)p_d}{m-k} - 2p_d}{\frac{dm}{kn} + \frac{dm}{(m-k)n}},$$

or

$$\frac{p_s - p_d}{pm},$$

which is the claimed expression. Since this holds for all $M \in I(n)$, the result follows. \blacksquare

Proof of Theorem 2: Note

$$\mathbf{Q}(\mathbf{P}, \mathbf{n}) = \frac{p_d}{p_s + (m-1)p_d} \mathbf{E}_m + \frac{p_s - p_d}{p_s + (m-1)p_d} \mathbf{I}_m,$$

where \mathbf{E}_m denotes the m -by- m matrix of ones and \mathbf{I}_m denotes the m -by- m identity matrix. Then, the eigenvalues of this matrix can be computed directly. The only nonzero eigenvalue of the first matrix is

$$\frac{mp_d}{p_s + (m-1)p_d}$$

with multiplicity 1; and adding

$$\frac{p_s - p_d}{p_s + (m-1)p_d} \mathbf{I}_m$$

just shifts all the eigenvalues by adding to them the constant multiplying the identity. Thus the second largest eigenvalue of $\mathbf{Q}(\mathbf{P}, \mathbf{n})$ (after the eigenvalue 1) is

$$\frac{p_s - p_d}{p_s + (m-1)p_d}.$$

Simple algebra shows that this is the same as the expression claimed in the theorem. \blacksquare

Proof of Corollary 4: Let us first show that with a probability going to 1 all nodes have degree between $(1 - f(n))d$ and $(1 + f(n))d$ for a function $f(n) \rightarrow 0$. From Lemma 2.1 in

Chung and Lu (2002),¹⁷ it follows that for any given i , $\mathbb{P}(d_i \geq (1 - f(n))d) > 1 - e^{-(f(n))^2 d/3}$ for any function $f(n) < 1$. The probability that all nodes have degrees at least $(1 - f(n))d$ is then at least $\left(1 - e^{-(f(n))^2 d/3}\right)^n$. Given that $d \geq \log^2(n)$, it follows that this expression is at least $\left(1 - \frac{e^{-(f(n))^2 \log(n)/3}}{n}\right)^n$, which goes to 1 as long as $(f(n))^2 \log(n)/3$ goes to ∞ . A similar argument establishes that all nodes have degrees at most $(1 + f(n))d$ with a probability going to 1. Thus, with a probability going to 1, $s_{\min} \geq \frac{1-f(n)}{(1+f(n))n}$.

Note, also that with a probability going to 1 that \mathbf{A} is connected (e.g., apply Theorem 1 noting that $h(n)$ is bounded away from 1 so that (iv) applies, and (i)-(iii) apply given the islands model and $d \geq \log^2(n)$). Thus, we can also apply Lemmas 3 and 4, to conclude that with a probability going to 1

$$\left\lfloor \frac{\log(n^2/4\gamma) - \log((1 + f(n))n/(1 - f(n)))}{-2 \log(|\lambda_2(\mathbf{T}(\mathbf{A}(\mathbf{P}, \mathbf{n})))|)} \right\rfloor \leq \text{CT}(\gamma/n^2; \mathbf{A}) \leq \left\lceil \frac{\log(n^2/\gamma)}{-2 \log(|\lambda_2(\mathbf{T}(\mathbf{A}(\mathbf{P}, \mathbf{n})))|)} \right\rceil,$$

and

$$\frac{\log(\frac{n}{2\gamma})}{-\log(|\lambda_2(\mathbf{T}(\mathbf{A}(\mathbf{P}, \mathbf{n})))|)} \leq \text{MT}(\gamma/n; \mathbf{A}(\mathbf{P}, \mathbf{n})) \leq \frac{\log(\frac{n}{2\gamma}) + \log((1 + f(n))n/(1 - f(n)))/2}{-\log(|\lambda_2(\mathbf{T}(\mathbf{A}(\mathbf{P}, \mathbf{n})))|)}.$$

These imply that with a probability going to 1:

$$\begin{aligned} \left\lfloor \frac{\log(n) - \log(4\gamma) - \log((1 + f(n))/(1 - f(n)))}{-2 \log(|\lambda_2(\mathbf{T}(\mathbf{A}(\mathbf{P}, \mathbf{n})))|)} \right\rfloor &\leq \text{CT}(\gamma/n^2; \mathbf{A}) & (10) \\ &\leq \left\lceil \frac{\log(n) - \log(\gamma)/2}{-\log(|\lambda_2(\mathbf{T}(\mathbf{A}(\mathbf{P}, \mathbf{n})))|)} \right\rceil, \end{aligned}$$

and

$$\begin{aligned} \frac{\log(n) - \log(2\gamma)}{-\log(|\lambda_2(\mathbf{T}(\mathbf{A}(\mathbf{P}, \mathbf{n})))|)} &\leq \text{MT}(\gamma/n; \mathbf{A}(\mathbf{P}, \mathbf{n})) & (11) \\ &\leq \frac{\frac{3}{2} \log(n) - \log(2\gamma) + \log((1 + f(n))/(1 - f(n)))/2}{-\log(|\lambda_2(\mathbf{T}(\mathbf{A}(\mathbf{P}, \mathbf{n})))|)}. \end{aligned}$$

Next, applying Theorem 2 and Lemma 5,

$$\left| \lambda_2(\mathbf{T}(\mathbf{A}(\mathbf{P}, \mathbf{n}))) - \frac{H(n) - 1}{m(n) - 1} \right| \xrightarrow{p} 0. \quad (12)$$

¹⁷Set the X_i 's in their lemma to be the realization of the links that a given node might have to other nodes.

Since, $H(n) = h(n)m(n)$ it follows that $\frac{H(n)-1}{m(n)-1} = \frac{h(n)m(n)-1}{m(n)-1}$ is bounded away from 1. Thus, from (12), we deduce that for any $1 > \delta > 0$, with a probability going to 1

$$\frac{1 - \delta}{-\log\left(\frac{H(n)-1}{m(n)-1}\right)} \leq \frac{1}{-\log(|\lambda_2(\mathbf{T}(\mathbf{A}(\mathbf{P}, \mathbf{n})))|)} \leq \frac{1 + \delta}{-\log\left(\frac{H(n)-1}{m(n)-1}\right)}.$$

The corollary then follows from (10) and (11), noting that $f(n) \rightarrow 0$. ▀

Proof of Theorem 5: For the left hand side, apply Theorem 3 and then compute the second eigenvalue of

$$\mathbf{Q}(\mathbf{P}) = \begin{bmatrix} \frac{f_1 p_s}{f_1 p_s + f_2 p_d} & \frac{f_2 p_d}{f_1 p_s + f_2 p_d} \\ \frac{f_1 p_d}{f_1 p_d + f_2 p_s} & \frac{f_2 p_s}{f_1 p_d + f_2 p_s} \end{bmatrix},$$

(the result appears in Jackson (2008), Section 8.3.6). For the right hand side, first use the definition of DWH; then apply Lemma 7 to show the degrees in the denominators in the DWH formula are arbitrarily close to their expectations; then use the strong law of large numbers to conclude that the ratios appearing in the formula converge to their expectations. ▀

7.3 Results and Proofs for the Empirical Analysis

7.3.1 Misidentifying Islands and Affine Bias

To justify including an intercept in our regressions, consider the following stylized elaboration of the islands model. We have m equally sized islands N_1, \dots, N_m and each of those islands k is divided into r equally sized sub-islands N_{k1}, \dots, N_{kr} . If i and j are in different islands, then the linking probability between them is p_d . If i and j are in the same island, then the linking probability depends on whether they are in the same sub-island, or different sub-islands. If they are in the same sub-island, then they are linked with probability p_s . And if they are in the same island but different sub-islands, then they are linked with probability p_b . We assume $0 < p_d \leq p_b \leq p_s$.

The idea is that the researcher has data on the islands but not the sub-islands. We will now study, in this simple setting, what happens if the homophily H is estimated as if the data were generated by the islands model with islands N_1, \dots, N_m and no sub-islands.

If, without knowing about the sub-islands, we estimate the probability of same-type nodes being linked, we are actually estimating the quantity

$$\tilde{p}_s = \frac{p_s + (r - 1)p_b}{r},$$

and our estimate of H , the unnormalized homophily, will be

$$\tilde{H} = \frac{\tilde{p}_s}{mp},$$

where p will be estimated correctly by its sample analogue of link density.

In this setting, it is not valid to apply Corollary 3 with the predictor of the second eigenvalue computed based on the misidentified island structure. That is, the second eigenvalue will not be well estimated by $\frac{\tilde{H}-1}{m-1}$. Instead, the second eigenvalue of the representative agent matrix¹⁸ will be an affine function of $\frac{\tilde{H}-1}{m-1}$. This is the content of the following proposition.

PROPOSITION 7. *In the modified islands setting just described, if $\tilde{x} = \frac{\tilde{H}-1}{m-1}$ is the regressor computed without information about the sub-island structure, then*

$$\lambda_2(\mathbf{Q}(\mathbf{P}, \mathbf{n})) = \alpha\tilde{x} + \beta,$$

where α and β depend on m and r .

Proof of Proposition 7: Letting \mathbf{E}_k denote the matrix of all ones of size k and \mathbf{I}_k denote the identity matrix of size k , we find that with \mathbf{P} specified by the description above,

$$\mathbf{Q}(\mathbf{P}, \mathbf{n}) = \frac{p_d\mathbf{E}_{mr} + (p_b - p_d)\mathbf{I}_m \otimes \mathbf{E}_r + (p_s - (p_b - p_d))\mathbf{I}_{mr}}{p_s + (r-1)p_b + (m-1)rp_d}.$$

Now, the second eigenvalue of

$$p_d\mathbf{E}_{mr} + (p_b - p_d)\mathbf{I}_m \otimes \mathbf{E}_r$$

is $r(p_b - p_d)$. Thus,

$$\lambda_2(\mathbf{Q}(\mathbf{P}, \mathbf{n})) = \frac{r(p_b - p_d) + p_s - (p_b - p_d)}{p_s + (r-1)p_b + (m-1)rp_d}.$$

This can be rewritten as

$$\lambda_2(\mathbf{Q}(\mathbf{P}, \mathbf{n})) = \frac{\tilde{H}m(mr-1) + 1 - r}{(m-1)r}.$$

Letting $\tilde{x} = \frac{\tilde{H}-1}{m-1}$ be the regressor, we find that

$$\lambda_2(\mathbf{Q}(\mathbf{P}, \mathbf{n})) = \frac{m(mr-1)}{r} \cdot \tilde{x} + \frac{mr+r-1}{r}.$$

¹⁸This, by Theorem 3 is the limit of the second eigenvalue of the realized matrix.

Now, in running the regressions, we do not make use of the details of the formula in the proof. We merely note that there is an affine bias if there is some homophily inside the islands on dimensions falling outside the scope of our data. Thus, including an intercept in the regression of convergence rates on \tilde{x} is a reasonable first-order approximation to account for some of this affine bias. ■

Of course, in more realistic settings, the various kinds of symmetry present in this model will not exist. However, it appears that more general formulas or characterizations could be obtained describing how homophilies at various levels interact. This could be a useful direction to pursue in taking this model to empirical settings, where there will almost always be some underlying homophily on dimensions not captured by the data.

7.3.2 The Asymptotic Equivalence of DWH and EDWH

Proof of Lemma 6: Consider the islands model with $m(n) \geq 2$ equal-sized groups, and probabilities of links within and across types p_s and p_d , respectively, and consider any sequence of groupings of islands $M(n) \in I(n)$. We show that

$$|\text{EDWH}(M(n), m, p_s, p_d) - \text{DWH}(M(n), \mathbf{A}(\mathbf{P}, \mathbf{n}))| \xrightarrow{p} 0.$$

This follows from showing that $|EW_{M(n), M(n)} - W_{M(n)M(n)}(\mathbf{A}(\mathbf{P}, \mathbf{n}))| \xrightarrow{p} 0$, $|EW_{M(n), M^c(n)} - W_{M(n)M^c(n)}(\mathbf{A}(\mathbf{P}, \mathbf{n}))| \xrightarrow{p} 0$, and $(\sum_{i \in M(n)} \frac{1}{d(n)}) / (\frac{1}{d_i(\mathbf{A}(\mathbf{P}, \mathbf{n}))}) \xrightarrow{p} 1$, for any $M(n)$ (and noting that the denominator is bounded away from 0 in the limit).

The latter conclusion follows from the argument in the proof of Corollary 4 using Lemma 2.1 in Chung and Lu (2002) to show that with a probability going to 1 all nodes have degree between $(1 - f(n))d(n)$ and $(1 + f(n))d(n)$ for any function $f(n) \rightarrow 0$ such that $(f(n))^2 \log(n)/3$ goes to ∞ . Next, given that $m(n)/n \rightarrow 0$ which implies that the number of nodes in within any island $i(n) = n/m(n)$ is growing without bound, and so we can again apply Lemma 2.1 in Chung and Lu (2002),¹⁹ to deduce that there is a function $g(n) \rightarrow 0$ such that with a probability going to 1,

$$\begin{aligned} (1 - g(n)) [k(n)p_s(n)|i(n)|^2/d(n)^2 + p_d(n)k(n)(k(n) - 1)|i(n)|^2/d(n)^2] &\leq \sum_{i \in M, j \in M} T_{ij}T_{ji} \\ &\leq (1 + g(n)) [k(n)p_s(n)|i(n)|^2/d(n)^2 + p_d(n)k(n)(k(n) - 1)|i(n)|^2/d(n)^2], \end{aligned}$$

¹⁹Now, we work with the X_i 's in their lemma to be the realization of the links within a given $M(n)$.

where $k(n)$ is the number of islands in $M(n)$ and similarly

$$(1-g(n))p_d(n)|M(n)||M^c(n)|/d(n)^2 \leq \sum_{i \in M, j \in M} T_{ij}T_{ji} \leq (1+g(n))p_d(n)|M(n)||M^c(n)|/d(n)^2.$$

These imply that

$$|EW_{M(n),M(n)} - W_{M(n)M(n)}(\mathbf{A}(\mathbf{P}, \mathbf{n}))| \xrightarrow{p} 0$$

and

$$|EW_{M(n),M^c(n)} - W_{M(n)M^c(n)}(\mathbf{A}(\mathbf{P}, \mathbf{n}))| \xrightarrow{p} 0,$$

as claimed. ■