Centrality in Stochastic Networks

Mohamed Mostagir*

James Siderius†

Abstract

Centrality measures are ubiquitous, appearing in models of opinion formation, macroeconomics, and consumption with externalities. With few exceptions, most of the previous literature has focused on modeling centrality in settings where the underlying network structure is known and remains static. This paper expands on this work by considering arbitrary row-stochastic random networks that may be evolving over time. Under mild assumptions, we show that all centrality measures are, with high probability, close to their values in an appropriately-defined "average" network. We conclude by demonstrating how this result offers a major technical simplification for the dynamic and stochastic analyses of several applications.

1 Introduction

Centrality is a fundamental concept in network analysis that captures not only the *number* of connections a particular entity has, but also the *significance* of these connections (see Jackson (2010)). An entity that forms connections with more central entities becomes more central herself. Many applications leverage network centrality for predictive power and decision making in a diverse range of domains: Ahern (2012) shows that centrality in intersectoral trade networks can be used as a predictor for stock performance; Golub and Jackson (2010) show that in models of naive learning in social networks, the influence of an agent on the consensus belief is directly linked to her eigenvector centrality in the normalized network; Candogan et al. (2012) show that optimal product pricing in a network with externalities should depend on agents' Bonacich centrality (Bonacich (1987)), and similarly, Acemoglu et al. (2015) show that the extent of contagion in the financial sector due to a liquidity shock at a certain bank is again closely related to the Bonacich centrality of that bank in the normalized financial network.

Contribution The vast majority of work on centrality assumes that the network structure is known and fixed. However, network data is noisy and networks are rarely static configurations:

*University of Michigan Ross School of Business

†Laboratory for Information and Decision Systems, MIT

people break ties and form new relationships over time, institutions dissolve partnerships and enter into new ones, etc. The question then is how to capture centrality in these dynamic and stochastic environments. We study a centrality measure based on the *normalized adjacency matrix* of the graph, which we call *row-stochastic centrality*. This measure covers a wide array of applications, including all the ones mentioned in the previous paragraph. Our main contribution is a result that shows that under mild regularity conditions, the centrality measures in these stochastic networks are with high probability close to their values in an appropriately-defined average network. This provides a methodological tool that, as we show in Section 4, greatly simplifies the study and analyses of these stochastic networks.

Related Literature While the literature mentioned earlier focuses on static and deterministic networks, there is some recent work on centrality in random networks. Using the graph-limit concept of graphons (see Parise and Ozdaglar (2019) for a primer), Avella-Medina et al. (2020) provide concentration inequalities for many common centrality measures, such as eigenvector and Katz-Bonacich centrality. For discrete graphs, Dasaratha (2017) shows that these centrality measures converge to their expected values with high probability. Our work complements these papers by considering a different centrality measure that captures a wide range of applications (e.g. all the aforementioned applications) for which the above results do not apply. This is because these applications use the normalized adjacency matrix instead of the adjacency matrix itself, which means that while the matrix is now row-stochastic, it is no longer symmetric. This breakage of symmetry presents additional technical challenges beyond the proof techniques used in these papers.

1.1 Deterministic Network Model

We now introduce our centrality measure and provide some simple examples where this type of centrality arises endogenously as a key measure of interest.

We consider an unweighted, undirected network $\mathbf{G}^*$ on n nodes with symmetric adjacency matrix $\mathbf{A}^*$ with elements a^*_{ij}. The corresponding weighted, directed network $\mathbf{G}$ is formed by normalizing $\mathbf{A}^*_i$ by agent i's degree to form a row-stochastic adjacency matrix $\mathbf{A}$ with elements a_{ij}. Precisely, $d_i = \sum_{j=1}^n a^*_{ij}$ and $a_{ij} = a^*_{ij}/d_i$.

Every node i in the network has a type $\gamma_i \in [0, 1]$, and a propensity $\theta_i \in [0, 1]$ corresponding to how "anchored" she is to her type. The *row-stochastic centrality* C_i of agent i is a solution to the fixed-point problem:

$$C_i = \theta_i \gamma_i + (1-\theta_i)\sum_{j=1}^{n} a_{ij} C_j \tag{1}$$

It can be easily shown that if $\mathbf{G}^*$ is connected and $\theta_i > 0$ for some agent i, the centralities have a unique solution, with the closed-form representation $C_i = \sum_{k=0}^{\infty} \mathbf{A}_\theta^k (\boldsymbol{\theta} \odot \boldsymbol{\gamma})$ where:[1]

$$\mathbf{A}_\theta = \begin{pmatrix} 1-\theta_1 & 0 & \cdots & 0 \\ 0 & 1-\theta_2 & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & \cdots & 1-\theta_n \end{pmatrix} \mathbf{A}$$

While $\mathbf{A}$ is a row-stochastic matrix, $\mathbf{A}_\theta$ is a row-substochastic and irreducible matrix, so the above sum always exists. As mentioned, this measure captures many applications. For example:

Opinion Formation. Consider a model of opinion formation where agents are either "stubborn" or "normal" and agent i holds opinion $\pi_{i,t} \in [0,1]$ over discrete time t, as in Yildiz et al. (2013). If an agent is stubborn, she always holds an opinion that is either $\pi_{i,t} = 1$ or $\pi_{i,t} = 0$, which is given exogenously and unchanged after $t = 0$. If the agent is normal, she begins with some initial opinion $\pi_{i,0}$ and updates this opinion by linearly combining the opinions of her neighbors $N(i) \subset \{1, \ldots, n\}$:

$$\pi_{i,t+1} = \frac{1}{|N(i)|} \sum_{j \in N(i)} \pi_{j,t}$$

The goal is to understand the limiting opinion of agents as a function of network structure and agent types. This problem can be reduced to our normalized centrality measure where we assign $\gamma_i = \pi_{i,0}$ and let $\theta_i = 1$ for stubborn agents (i.e. they are completely anchored to their type) and $\theta_i = 0$ otherwise. The centrality measure then captures exactly the opinion of agent i as $t \to \infty$.

Contagion and Input-Output Economies. Consider an economy consisting of n different sectors with Cobb-Douglas production functions (see Burres (1985)). Sector i has an exogenous demand $D_i > 0$ for its products, but sectors may also require other sectors' outputs as inputs to their own production. We can write down P_i, the production of sector i as:

$$P_i = D_i + \sum_{j=1}^{n} \omega_{ij} P_j$$

for some shares $\{\omega_{ij}\}_{i,j}$, where ω_{ij} denotes the fraction of sector i's output that sector j uses. In a similar vein to Acemoglu et al. (2011), we assume that $\sum_{j=1}^{n} \omega_{ij} < 1$, so that there are (at most)

[1] See Mostagir et al. (2019) for a proof.

constant returns to scale.

This model provides a framework for the interconnections of the economy and the corresponding effects on output. A shock to the demand (or productivity) at sector i (so that $D_i' < D_i$) affects not only the production of that sector but potentially also the production of some sector j. This indirect impact is a result of spillover contagion, and can be measured precisely through agent j's centrality (and its change after the demand shock at i).

When the input-output network has every sector contributing equally to those it supplies, we can cast the output of sector i in the form of our centrality measure. Let us first normalize the equation by constant $\bar{D} = \max_i(1 - \sum_j \omega_{ij})^{-1} D_i$:

$$\frac{P_i}{\bar{D}} = \frac{D_i}{\bar{D}} + \sum_{j=1}^{n} \omega_{ij} \frac{P_j}{\bar{D}}$$

Then setting $\theta_i = 1 - \sum_j \omega_{ij}$, $\gamma_i = (1 - \sum_j \omega_{ij})^{-1} D_i / \bar{D}$, and $a_{ij} = (\sum_j \omega_{ij})^{-1} \omega_{ij}$, we obtain the reduction.

Consumption with Externalities. Let us consider the model of Candogan et al. (2012), whereby agents' utilities from consumption depend on other agents' consumption decisions. In particular:

$$u_i = \alpha_i x_i - b_i x_i^2 + x_i \cdot \sum_{j=1}^{n} g_{ij} x_j - p_i x_i$$

where α_i, b_i are constants, p_i is the price for agent i, and g_{ij} terms capture the externalities from agent j on agent i. For a given set of prices $\mathbf{p}$, the equilibrium consumption is given recursively by:

$$x_i = \frac{\alpha_i - p_i}{2b_i} + \frac{1}{2b_i} \sum_{j=1}^{n} g_{ij} x_j$$

By Assumption 1 in Candogan et al. (2012), we have that $\sum_{j=1}^{n} g_{ij} < 2b_i$, so in particular letting $\theta_i = 1 - \sum_{j=1}^{n} g_{ij}/2b_i$, we see that $\theta_i \in (0,1)$. Finally, setting $a_{ij} = \frac{g_{ij}}{2b_i(1-\theta_i)}$ and $\gamma_i = \frac{\alpha_i - p_i}{2b_i \theta_i}$,[2] we get the reduction to our centrality measure in Equation (1).

[2]Note that $\gamma_i > 0$ when $\alpha_i > p_i$ (marginal utility at zero consumption exceeds price). On the other hand, it may be possible that $\gamma_i > 1$; simply note that dividing by a sufficiently large constant c on each side, we can always recast the problem with $\gamma_i \in [0,1]$ using x_i/c instead of x_i.

2 Random Network Model

We start by defining some notation for a random network model on n agents. This random network is specified by a matrix $\boldsymbol{\rho}$ of link probabilities, where:

$$\boldsymbol{\rho} = \begin{pmatrix} \rho_{11} & \rho_{12} & \cdots & \rho_{1n} \\ \rho_{21} & \rho_{22} & \cdots & \rho_{2n} \\ \cdots & \cdots & \cdots & \cdots \\ \rho_{n1} & \rho_{n2} & \cdots & \rho_{nn} \end{pmatrix}$$

with the probability that agent i is linked to agent j equal to ρ_{ij}. We let $a^*_{ij} \in \{0,1\}$ denote whether a link from i to j is realized. Take χ to be any sequence of realized links $(a^*_{i_1 j_1}, a^*_{i_2 j_2}, \cdots, a^*_{i_k j_k})$, and moreover let $\boldsymbol{\chi}_{-ij}$ denote the set of χ where both a^*_{ij} and a^*_{ji} do not appear. For additional simplification, we impose:

Assumption 1. For every i, j, the conditional probability of link formation satisfies $\mathbb{P}[a^*_{ij} = 1|\chi] = \rho_{ij}$ and $\mathbb{P}[a^*_{ij} = a^*_{ji}] = 1$, for all $\chi \in \boldsymbol{\chi}_{-ij}$.

In other words, (i) the probability of a link $i_1 \rightleftarrows j_1$ forming does not depend on whether $i_2 \rightleftarrows j_2$ forms, unless $i_1 = i_2$ *and* $j_1 = j_2$, and (ii) the link $i \to j$ exists if and only if $j \to i$ also exists (and hence we write $i \rightleftarrows j$). This defines an unweighted, undirected network $\mathbf{G}^*$ whose adjacency matrix is symmetric.

We consider two (weighted) network objects: (i) the realized network $\tilde{\mathbf{G}}$, and (ii) the "average" network $\bar{\mathbf{G}}$. In the realized network $\tilde{\mathbf{G}}$, we have weights given by:

$$\tilde{a}_{ij} = \begin{cases} (1-\theta_i)/d_i, & \text{if } a^*_{ij} = 1 \\ 0, & \text{otherwise} \end{cases}$$

where $d_i = \sum_{j=1}^n a^*_{ij}$ is the realized degree of agent i. If $d_i = 0$, then we set $\tilde{a}_{ii} = 1 - \theta_i$ and $\tilde{a}_{ij} = 0$ for all $i \neq j$. On the other hand, in the expected network $\bar{\mathbf{G}}$, expected weights are given by:

$$\bar{a}_{ij} = (1-\theta_i)\rho_{ij}/\bar{d}_i$$

where $\bar{d}_i = \sum_{j=1}^n \rho_{ij}$, which is *expected degree* of agent i.[3] As before, if $\bar{d}_i = 0$, then we set $\bar{a}_{ii} = 1 - \theta_i$ and $\bar{a}_{ij} = 0$ for all $i \neq j$. Note that here that the corresponding matrices $\tilde{\mathbf{A}}$ and $\bar{\mathbf{A}}$ are in the row-substochastic representation instead of the row-stochastic representation.

[3] Note that $\bar{a}_{ij}$ is not technically the expectation of $\tilde{a}_{ij}$ for finite n, but these two expressions are shown to be consistent as $n \to \infty$.

Dynamic Networks. Consider a stochastic sequence of realized networks $\tilde{\mathbf{G}}^{(t)}$ defined on n agents over discrete time $t = 1, 2, \ldots$. Let $\mathcal{F}^{(t)}$ be the filtration $\mathcal{F}^{(t)} = \sigma(\tilde{\mathbf{G}}^{(\tau)} | \tau < t)$ and $\boldsymbol{\rho}^{(t)}$ be the random network distribution at time t with respect to the filtration $\mathbb{F} \equiv \{\mathcal{F}^{(t)}\}_{t=1}^{n}$. Then, provided that $\boldsymbol{\rho}^{(t)}$ satisfies Assumption 1 with probability 1 for all t, our random network analysis applies identically to the case of dynamic networks.

Main Result. Without loss, let us focus on the case of a random (but static) network. Consider a sequence of growing societies $\mathcal{S}_n$ each with n agents. The main connection we develop in this section is between a growing sequence of random networks $\tilde{\mathbf{G}}_n$ and the corresponding "average" networks given by $\bar{\mathbf{G}}_n$. Agent i in society $\mathcal{S}_n$ has weight $\theta_i^{(n)} \in (0, 1)$,[4] and for simplicity, we write $\boldsymbol{\theta}^{(n)}$ as the vector of θ's when the population is of size n. We also have the expected degree matrix $\bar{\mathbf{D}}_n$ given by:

$$\bar{\mathbf{D}}_n = \begin{pmatrix} \sum_{j=1}^{n} \rho_{1j} & 0 & \cdots & 0 \\ 0 & \sum_{j=1}^{n} \rho_{2j} & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & \cdots & \sum_{j=1}^{n} \rho_{nj} \end{pmatrix}$$

Finally, we make the following assumption:

Assumption 2. Consider the sequence of vectors $\boldsymbol{\theta} \equiv \{\boldsymbol{\theta}^{(n)}\}_{n=1}^{\infty}$ along with the expected degree matrix $\bar{\mathbf{D}}_n$ and the link probability matrix $\boldsymbol{\rho}_n$. Then (i) $\min_i \theta_i^{(n)} > 0$ for all n; and (ii) *Laplacian matrix*, $\bar{\mathbf{D}}_n - \boldsymbol{\rho}_n$, has its second-smallest eigenvalue bounded away from 0.

The first condition requires that *some* agent incorporates their type directly into their centrality. The second condition requires the Laplacian matrix of the realized network not have multiple zero eigenvalues, as this would imply the network is not connected. Together, these conditions ensure the centrality measure is well-defined.

Each society $\mathcal{S}_n$ comes with its own random network generation process given by:

$$\boldsymbol{\rho}_n = \begin{pmatrix} \rho_{11}(n) & \rho_{12}(n) & \ldots & \rho_{1n}(n) \\ \rho_{21}(n) & \rho_{22}(n) & \ldots & \rho_{2n}(n) \\ \cdots & \cdots & \cdots & \cdots \\ \rho_{n1}(n) & \rho_{n2}(n) & \ldots & \rho_{nn}(n) \end{pmatrix}$$

We present the main result of our paper, which includes two regularity conditions similar to those in Dasaratha (2017). Our first condition requires that agents' degrees grow at a sufficiently

[4] We remark that under some condition, this does not automatically rule out applications where θ can be equal to 0 or 1, as we discuss after Definition 2 and in Footnote 6.

fast rate. This is important because in very sparse networks, individual realizations of links have a significant effect on centrality. For example, when the expected degree is bounded above as $n \to \infty$, then any individual link represents a non-vanishing contribution toward an agent's centrality.

Definition 1 (Expected Degrees)**.** We say that random network generation $\boldsymbol{\rho}_n$ satisfies the *expected-degrees* condition if $\lim_{n\to\infty} \min_{i\in\mathcal{S}_n} \bar{d}_i^{(n)} / \log n = \infty$.

In other words, the expected degrees condition requires that as the society $\mathcal{S}_n$ grows, *all* agents in the society have expected degrees which are uniformly growing with $\log n$, regardless of the time of their birth.[5] Note that this is a stronger condition than just requiring every agent in the network to have an expected degree that grows strictly faster than $\log n$. The difference between the two conditions may be subtle, but without the stronger condition, there is a significant chance that the realized network has row-stochastic centralities bounded away from those of the average network, even as the population size grows large. For the interested reader, this difference is illustrated in Example 1.

Secondly, recall from Equation (1) that agents split their centrality by assigning weights to their type and the types of their neighbors. The next definition states that society needs to be somewhat homogeneous, in the sense that agents have a "similar enough" distribution of weights:

Definition 2 (Normal Society)**.** We say that $\boldsymbol{\theta}$ satisfies the *normal society* condition if there exists a constant $\nu < \infty$ such that for all agents i, j:

$$\limsup_{n\to\infty} \frac{\theta_i^{(n)}}{\theta_j^{(n)}} \leq \nu$$

As a special case, the "normal society" condition is always satisfied when all agents have the exact same θ. Interested readers can refer to Example 2 to see how, in the absence of the normal society condition, two networks that occur with equal probabilities can have different properties as $n \to \infty$. It is easy to see that the normal society condition may also be relaxed so that it applies only to agents with $\gamma_i > 0$ (so, for instance, the stubborn opinion formation dynamics in Section 1 can still leverage Theorem 1 despite having $\theta = 0$ or 1).

Under the above regularity conditions, we get the following reduction to deterministic network analysis:

[5] Formally, Definition 1 can be replaced with the "uniformity requirement" that $\exists \beta_n$ such that $d_i^{(n)} \geq \beta_n \log n$ for all $i \leq n$ with $\lim_{n\to\infty} \beta_n = \infty$.

Theorem 1. *Suppose Assumptions 1 and 2 hold, the sequence of $\boldsymbol{\rho}_n$ satisfies the expected-degree condition, and $\boldsymbol{\theta}$ is a normal society. Then, the centrality vector $\tilde{\mathbf{C}}^{(n)}$ in random network $\tilde{\mathbf{G}}_n$ is with high probability close to its centrality vector $\bar{\mathbf{C}}^{(n)}$ in the average network $\mathbf{G}_n$. In other words, for all $\epsilon > 0$,*

$$\lim_{n\to\infty} \mathbb{P}\left[\left|\left|\tilde{\mathbf{C}}^{(n)} - \bar{\mathbf{C}}^{(n)}\right|\right|_{\infty} > \epsilon\right] = 0$$

Proof. Let us denote by $\bar{\mathbf{E}}_n$ the "expected" (normalized) adjacency matrix, $\bar{\mathbf{E}}_n = \mathbb{E}[\tilde{\boldsymbol{\rho}}_n]\mathbb{E}[\tilde{\mathbf{D}}_n]^{-1}$, and the "mean" influence network, $\bar{\mathbf{A}}_n = \mathbf{O}_{\theta}^{(n)}\bar{\mathbf{E}}_n$, where:

$$\mathbf{O}_{\theta}^{(n)} = \begin{pmatrix} (1-\theta_1^{(n)}) & 0 & \dots & 0 \\ 0 & (1-\theta_2^{(n)}) & \dots & 0 \\ \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & (1-\theta_n^{(n)}) \end{pmatrix}$$

The first step of the proof establishes that the difference between the "mean" (normalized) adjacency matrix and the realized (normalized) adjacency matrix, $||\tilde{\boldsymbol{\rho}}_n\tilde{\mathbf{D}}_n^{-1} - \bar{\mathbf{E}}_n||_2$, is small with high probability. In the second step, we prove that for any sequence of types $\boldsymbol{\gamma}_n$, the difference between the expected and realized centrality vector, $||\tilde{\mathbf{C}}(\boldsymbol{\gamma}_n) - \bar{\mathbf{C}}(\boldsymbol{\gamma}_n)]||_{\infty}$, is also small with high probability. Finally, we combine these two steps with Assumption 2 to show that the network is connected with high probability, so all measures are well-defined, completing the proof.

Step 1: Call $\tilde{\mathbf{E}}_n \equiv \tilde{\boldsymbol{\rho}}_n\tilde{\mathbf{D}}_n^{-1}$. Let $\psi > 0$. Let $\underline{d}^{(n)} = \min_i d_i^{(n)}$; that is, $\underline{d}^{(n)}$ is the expected minimum degree. We first show that the Laplacian matrices $\tilde{\mathcal{L}}_n = \mathbf{I} - \tilde{\mathbf{D}}_n^{-1/2}\tilde{\boldsymbol{\rho}}_n\tilde{\mathbf{D}}_n^{-1/2}$ and $\bar{\mathcal{L}}_n = \mathbf{I} - \bar{\mathbf{D}}_n^{-1/2}\bar{\boldsymbol{\rho}}_n\bar{\mathbf{D}}_n^{-1/2}$ satisfy $\lim_{n\to\infty}\mathbb{P}[||\tilde{\mathcal{L}}_n - \bar{\mathcal{L}}_n||_2 \geq \psi] = 0$ (i.e., they are equal with high probability). It follows from Theorem 2 in Chung and Radcliffe (2011) that with probability at least $1-\psi$:

$$||\tilde{\mathcal{L}}_n - \bar{\mathcal{L}}_n||_2 \leq 2\sqrt{\frac{3\log(4n/\psi)}{\underline{d}^{(n)}}}$$

By the expected-degrees condition, we know that $\lim_{n\to\infty}\underline{d}^{(n)}/\log n \to \infty$, which implies that:

$$\limsup_{n\to\infty} ||\tilde{\mathcal{L}}_n - \bar{\mathcal{L}}_n||_2 \leq \lim_{n\to\infty} 2\sqrt{\frac{3\log(4n/\psi)}{\underline{d}^{(n)}}} = 0$$

establishing the desired result. It is clear that the same implication is true of the matrices:

$$\tilde{\mathbf{N}} \equiv \tilde{\mathbf{D}}_n^{-1/2}\tilde{\boldsymbol{\rho}}_n\tilde{\mathbf{D}}_n^{-1/2} = \tilde{\mathbf{D}}_n^{-1/2}\tilde{\mathbf{E}}_n\tilde{\mathbf{D}}^{1/2}$$

$$\bar{\mathbf{N}} \equiv \bar{\mathbf{D}}_n^{-1/2}\bar{\boldsymbol{\rho}}_n\bar{\mathbf{D}}_n^{-1/2} = \bar{\mathbf{D}}^{-1/2}\bar{\mathbf{E}}_n\bar{\mathbf{D}}^{1/2}$$

Thus, let us write:

$$\begin{aligned}
\limsup_{n\to\infty} ||\tilde{\mathbf{E}}_n - \bar{\mathbf{E}}_n||_2 &\leq \limsup_{n\to\infty} ||\tilde{\mathbf{D}}^{1/2}\tilde{\mathbf{N}}\tilde{\mathbf{D}}^{-1/2} - \bar{\mathbf{D}}^{1/2}\bar{\mathbf{N}}\bar{\mathbf{D}}^{-1/2}||_2 \\
&\leq \limsup_{n\to\infty} \left(\max\{||\tilde{\mathbf{D}}^{1/2}||_2, ||\bar{\mathbf{D}}^{1/2}||_2\}\right) ||\tilde{\mathbf{N}} - \bar{\mathbf{N}}||_2 \left(\max\{||\tilde{\mathbf{D}}^{-1/2}||_2, ||\bar{\mathbf{D}}^{-1/2}||_2\}\right) \\
&\leq \limsup_{n\to\infty} \psi \cdot \left(\max\{||\tilde{\mathbf{D}}^{1/2}||_2, ||\bar{\mathbf{D}}^{1/2}||_2\}\right) \cdot \left(\max\{||\tilde{\mathbf{D}}^{-1/2}||_2, ||\bar{\mathbf{D}}^{-1/2}||_2\}\right)
\end{aligned}$$

Recall that a^*_{ij} is the binary random variable for whether there exists a link $i \to j$; note that:

$$\begin{aligned}
||\tilde{\mathbf{D}}_n^{1/2}||_2||\bar{\mathbf{D}}_n^{-1/2}||_2 &= \sqrt{\frac{\max_i \sum_{j=1}^n a^*_{ij}}{\max_i \sum_{j=1}^n \rho_{ij}}} \leq \max_i \sqrt{\frac{\sum_{j=1}^n a^*_{ij}}{\sum_{j=1}^n \rho_{ij}}} \\
||\tilde{\mathbf{D}}_n^{1/2}||_2||\bar{\mathbf{D}}_n^{-1/2}||_2 &= \sqrt{\frac{\max_i \sum_{j=1}^n \rho_{ij}}{\max_i \sum_{j=1}^n a^*_{ij}}} \leq \max_i \sqrt{\frac{\sum_{j=1}^n \rho_{ij}}{\sum_{j=1}^n a^*_{ij}}}
\end{aligned}$$

which are both bounded above almost surely. To see this, note that for n large, we can apply the Lyapunov Central Limit Theorem (see Billingsley (1995)):

$$\begin{aligned}
\lim_{n\to\infty} \frac{\sum_{j=1} a^*_{ij}}{\sum_{j=1}^n \rho_{ij}} - 1 &\sim \frac{1}{\sum_{j=1}^n \rho_{ij}} N\left(0, \sum_{j=1}^n \rho_{ij}(1-\rho_{ij})\right) \\
&\sim \frac{1}{\sqrt{\log n}} N\left(0, \Omega_n\right)
\end{aligned}$$

where $\Omega_n \to 0$. If $z_1, \ldots, z_n$ are normally distributed with variance σ^2, then by the Fisher-Tippet-Gnedenko theorem (see Charras-Garrido and Lezaud (2013) and Taylor (2011)), we see that:

$$\mathbb{E}\left[\max_i z_i\right] \in O(\sigma\sqrt{\log n})$$

Therefore, we have by Jensen's inequality:

$$\begin{aligned}
\mathbb{E}\left[\max_i \sqrt{\frac{\sum_{j=1}^n a^*_{ij}}{\sum_{j=1}^n \rho_{ij}}}\right] &= \mathbb{E}\left[\sqrt{\max_i \frac{\sum_{j=1}^n a^*_{ij}}{\sum_{j=1}^n \rho_{ij}}}\right] \\
&\leq \sqrt{\mathbb{E}\left[\max_i \frac{\sum_{j=1}^n a^*_{ij}}{\sum_{j=1}^n \rho_{ij}}\right]} \in 1 + O(\sqrt{\Omega_n}) \overset{n\to\infty}{\to} 1
\end{aligned}$$

which by Markov's inequality suggests for any $\kappa > 0$:

$$\lim_{n\to\infty} \mathbb{P}\left[\max_i \sqrt{\frac{\sum_{j=1}^n a^*_{ij}}{\sum_{j=1}^n \rho_{ij}}} \leq (1+\kappa)\right] = 1$$

Similar reasoning proves the second case where $\max_i \tilde{\mathbf{D}}_n < \max_i \bar{\mathbf{D}}_n$. This establishes that $||\tilde{\mathbf{E}}_n -$

$\bar{\mathbf{E}}_n||_2$ is small with high probability.

Step 2: We note then that $\tilde{\mathbf{A}}_n = \mathbf{O}_\theta^{(n)}\tilde{\mathbf{E}}_n$ and $\bar{\mathbf{A}}_n = \mathbf{O}_\theta^{(n)}\bar{\mathbf{E}}_n$. Fix a sequence of influence vectors, $\boldsymbol{\gamma}_n$. For every $\psi > 0$, we can write for large enough n:

$$\begin{aligned}
\left|\left|(\mathbf{I}-\tilde{\mathbf{A}}_n)^{-1}-(\mathbf{I}-\bar{\mathbf{A}}_n)^{-1}\right|\right|_2 &= \left|\left|\sum_{k=0}^{\infty}\left(\tilde{\mathbf{A}}_n^k-\bar{\mathbf{A}}_n^k\right)\right|\right|_2 \\
&= \left|\left|\sum_{k=0}^{\infty}\mathbf{O}_n^k(\tilde{\mathbf{E}}_n^k-\bar{\mathbf{E}}_n^k)\right|\right|_2 \\
&\leq \sum_{k=0}^{\infty}\left|\left|\mathbf{O}_n^k\right|\right|_2\left|\left|\tilde{\mathbf{E}}_n^k-\bar{\mathbf{E}}_n^k\right|\right|_2 \\
&\leq \sum_{k=0}^{\infty}\left(1-\inf_i\theta_i^{(n)}\right)^k\psi \\
&\leq \frac{\psi}{\inf_i\theta_i^{(n)}}
\end{aligned}$$

Note that this implies for large n and any $\boldsymbol{\gamma}_n$:

$$\begin{aligned}
\left((\mathbf{I}-\tilde{\mathbf{A}}_n)^{-1}-(\mathbf{I}-\bar{\mathbf{A}}_n)^{-1}\right)\left(\boldsymbol{\gamma}_n\odot\boldsymbol{\theta}^{(n)}\right) &\leq \frac{\psi}{\inf_i\theta_i^{(n)}}\mathbf{1}\odot\boldsymbol{\theta}^{(n)} \\
&\leq \frac{\psi\sup_i\theta_i^{(n)}}{\inf_i\theta_i^{(n)}}\mathbf{1}\leq\psi\nu\mathbf{1}
\end{aligned}$$

by the normal society condition on $\boldsymbol{\theta}^{(n)}$.[6] Therefore, we can bound this difference from above by any constant $\zeta > 0$; in particular, for any $\zeta > 0$:

$$\lim_{n\to\infty}\mathbb{P}\left[\left|\left|\left((\mathbf{I}-\tilde{\mathbf{A}}_n)^{-1}-(\mathbf{I}-\bar{\mathbf{A}}_n)^{-1}\right)\left(\boldsymbol{\gamma}\odot\boldsymbol{\theta}^{(n)}\right)\right|\right|_\infty\geq\zeta\right]=0$$

Thus, as $n\to\infty$, with high probability we have for every $\mu > 0$ and $\boldsymbol{\gamma}_n$:

$$\lim_{n\to\infty}\mathbb{P}\left[\left|\left|\tilde{\mathbf{C}}(\boldsymbol{\gamma}_n)-\bar{\mathbf{C}}(\boldsymbol{\gamma}_n)]\right|\right|_\infty\geq\mu\right]=0$$

as desired.

Step 3: Under Assumption 2, $\bar{\rho}_n\bar{\mathbf{D}}_n^{-1}$ has a non-vanishing spectral gap, then for sufficiently large n, we know any two nodes i and j in the realized network $\tilde{\mathbf{G}}_n$ are connected with high probability. To see this, we first construct a directed network T by assigning weights $t_{ij} = [\bar{\boldsymbol{\rho}}_n\bar{\mathbf{D}}_n^{-1}]_{ij}$.

[6] One should note that we do not require the normal society condition to hold for agent i if $\gamma_i = 0$.

The (normalized) Laplacian matrix of the directed network T is given by:

$$\mathcal{L} = \mathbf{I} - \bar{\mathbf{D}}_n^{-1/2} \bar{\boldsymbol{\rho}}_n \bar{\mathbf{D}}_n^{-1/2}$$

We use the result of Chung (2005) which provides a Cheeger inequality for directed networks: namely, that the conductance of T, $\phi(T)$, is bounded below by $\lambda_2^T/2$. Note that $\phi(T) = \phi(\bar{\mathbf{E}}_n)$, and therefore we know that as $n \to \infty$, $\phi(\bar{\mathbf{E}}_n) \geq (1-\eta)/2 \equiv \kappa > 0$, where η is ratio of the first two eigenvalues.

Consider the network $\bar{\mathbf{E}}_n^* = \bar{\mathbf{E}}_n \bar{\mathbf{D}}_n \left(\bar{\mathbf{D}}_n^*\right)^{-1}$ Then $\bar{\mathbf{E}}_n^*$ is symmetric, so by Chung and Radcliffe (2011) and the same reasoning as in Step 1 (along with the fact that $\bar{\mathbf{E}}_n^*$ and $\mathbf{I} - \bar{\mathcal{L}}_n^*$ have the same eigenvalues), we know that:

$$||\tilde{\lambda}_\mu^* - \bar{\lambda}_\mu^*||_2 \leq 2\sqrt{\frac{3\log(4n/\psi)}{\underline{d}^{(n)}}}$$

for $\mu = 1, 2$. Note that $\lim_{n\to\infty} \underline{d}^{(n)}/\log n = \infty$ since $\bar{\mathbf{E}}_n$ has a non-vanishing spectral gap (and the previous conductance argument), and $\lim_{n\to\infty} \underline{d}^{(n)}/\log n = \infty$ by the expected-degrees condition. This implies that with high probability, $\tilde{\mathbf{E}}_n^*$ has no vanishing spectral gap, and so using the standard Cheeger inequality proves that it is connected w.h.p. Since $\tilde{\mathbf{E}}_n^*$ is connected if and only if $\tilde{\mathbf{E}}_n$ is, we see that $\tilde{\mathbf{E}}_n$ is connected w.h.p. and therefore, the network is connected. Combining this with $\min_i \theta_i^{(n)} > 0$, we know that the centrality measure $\tilde{\mathbf{C}}$ is well-defined by Lemma 3 in Mostagir et al. (2019). $\square$

Theorem 1 states that under the conditions above, it suffices to consider $\bar{\mathbf{G}}_n$ instead of $\tilde{\mathbf{G}}_n$ as $n \to \infty$. This provides a major technical simplification in a variety of examples as we illustrate in Section 4.

3 Theory Failure Examples

In this section, we provide examples to show that the Expected Degrees and Normal Society conditions required for Theorem 1 cannot be dispensed with.

Example 1 (Non-Uniform Slow Degree Growth)**.** Suppose the society $\mathcal{S}_n$ has the link probability matrix:

$$\boldsymbol{\rho}_n = \begin{pmatrix} 0 & 1 & \cdots & 1 & 1/2 \\ 1 & 0 & \cdots & 1 & 0 \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ 1 & 1 & \cdots & 0 & 1 \\ 1/2 & 0 & \cdots & 1 & 0 \end{pmatrix}$$

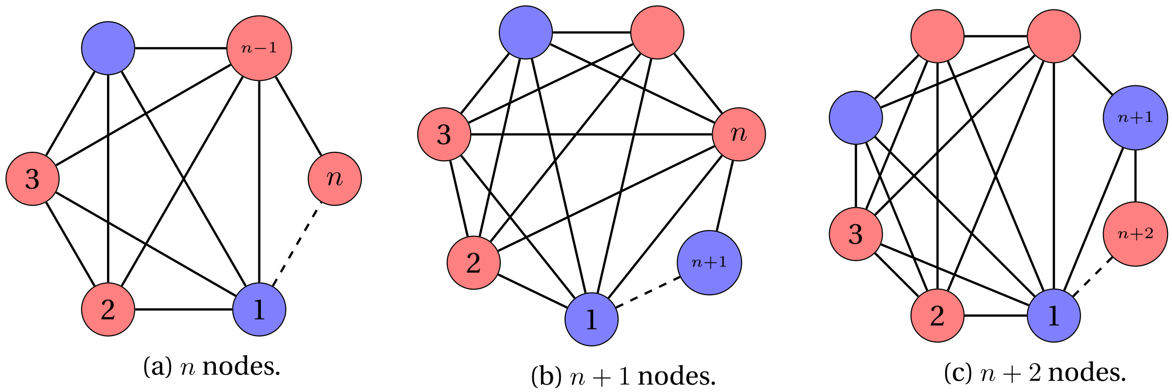

Figure 1. The expected network (solid links = weight 1, dashed lines = weight $1/2$).

In other words, the first $n-1$ nodes are arranged in a clique of size $n-1$, with the nth node adjacent node $n-1$ almost surely, and node n adjacent to node 1 with probability $1/2$. The expected network is given in Figure 1, whereas the possible realized networks are given in Figure 2. We assume that every two out of three (red) nodes have $\theta_i^{(n)} = 1/2$ and $\gamma_i = 1$, while one out of three (blue) nodes have $\theta_i^{(n)} = 1$ and $\gamma_i = 0$.

It can be shown via a walk-counting argument that the centrality of all red nodes $2 \leq i \leq (n-1)$ is at most approximately $2/3$ for large n. Similarly, for node n (in average network (a)), her centrality is equal to at most (approximately) $13/18 > 2/3$.

On the other hand, consider the realized network for society $\mathcal{S}_n$. In this case, $\tilde{\mathbf{G}}_n$ looks like one of the two networks in Figure 2, each with equal probability. For large n the centralities of all

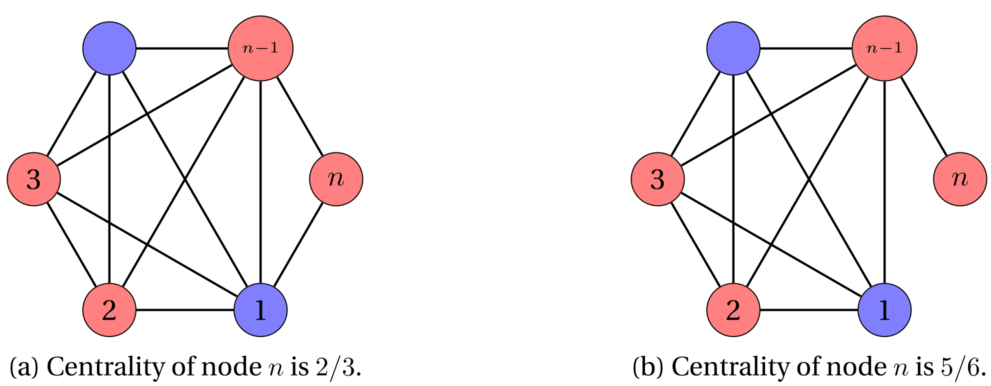

Figure 2. The realized network (all links are the same).

nodes $2 \leq i \leq (n-1)$ are approximately unchanged from the average network. In network (a), the maximal centrality of node n is strictly less than in the average network; in particular, it is

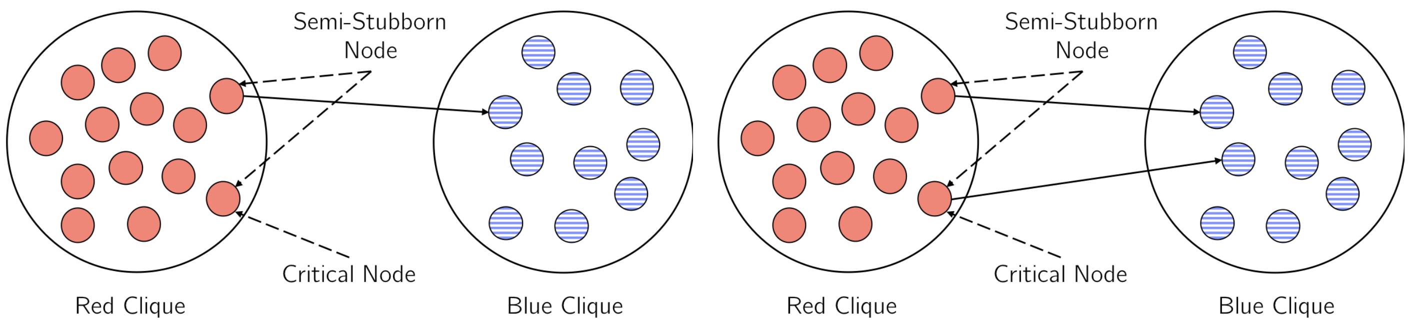

Figure 3. Example where the Non-normal Society assumption is violated.

equal to $2/3$, the same as all other nodes in the network. In network (b), the maximal centrality of node n is strictly more than in the average network; in particular, it is equal to $5/6$. Therefore, $\tilde{\mathbf{G}}_n$ in network (b) has a node with centrality $2/3$ with probability $1/2$ and centrality $5/6$ with probability $1/2$.

Of course, for any node i, as $n \to \infty$, the expected degree of node i grows faster than $\log n$ (it grows linearly!), but the minimum expected degree is constant. $\square$

Example 2 (Non-Normal Society)**.** For ease of notation, let us define $n_* = n - \log^2(n)$. Suppose the first $\log^2(n)$ nodes have $\theta_i = 1$ and $\gamma_i = 0$, the next two nodes have $\theta = 1/n_*$ and $\gamma_i = 1$ (who we will call *semi-stubborn*), and the remaining nodes have $\theta = 1/\exp(n_*)$ and $\gamma_i = 1$ (the former group is blue/shaded, and the latter two groups are red/solid, as seen in Figure 3.) All of the red nodes are pairwise adjacent with probability 1 in a clique; all of the blue nodes are pairwise adjacent with probability 1 in a clique as well. Suppose the first semi-stubborn red node is adjacent to the first blue node with probability 1 (so the network is always connected), whereas the second semi-stubborn node is adjacent to the first blue node with only probability $1/2$ (called the *critical node*). All other (non semi-stubborn) red nodes are never adjacent to a blue node.

When the critical red node is adjacent to the blue node, the walks to the blue nodes is given by:

$$w = \frac{2}{n_*} w_{ss} + \frac{n_* - 2}{n_*} w$$
$$w_{ss} = \frac{n_* - 1}{n_*^2} + \frac{(n_* - 2)(n_* - 1)}{n_*^2} w$$

which as $n \to \infty$ (so $n_* \to \infty$) satisfies $w = w_{ss} = 1/3$. Now, consider the case where the red node is *not* adjacent to the blue node; the walks to the blue nodes are given by:

$$\begin{aligned} w &= \frac{1}{n_*}w_c + \frac{1}{n_*}w_{ss} + \frac{n_*-2}{n_*}w \\ w_{ss} &= \frac{n_*-1}{n_*^2} + \frac{(n_*-1)(n_*-2)}{n_*^2}w + \frac{n_*-1}{n_*^2}w_c \\ w_c &= \frac{n_*-2}{n_*}w + \frac{1}{n_*}w_{ss} \end{aligned}$$

which as $n \to \infty$ satisfies $w = w_{ss} = w_c = 0$. Clearly the normal society condition is violated since the semi-stubborn nodes hold $\theta = 1/n_*$ and other red nodes hold $\theta = 1/\exp(n_*)$, and $\lim_{n\to\infty}(1/n_*)/(1/\exp(n_*)) = \infty$. And, consequently, the result of Theorem 1 cannot be applied.

4 Applications

Application 1: Opinion Dynamics with Friending/Unfriending. Consider a social network $\mathbf{G}^{(t)}$ which is evolving over time. Assume all agents are initially linked to a constant fraction of the population and $\mathbf{G}^{(t)}$ evolves according to the following process:

1. For every neighbor of agent i, delete an existing link $i \to j$ (i.e., unfriend) with probability p_i, and delete the symmetric link $j \to i$.

2. For every non-neighbor of agent i, add a non-existing link $i \to j$ (i.e., friend) with probability q_i, and add the symmetric link $j \to i$.

3. Run the stubborn-agent opinion dynamics and observe the change in limiting agent beliefs.

We are interested in understanding the opinions of agents in the network $\mathbf{G}^{(t)}$, given the network looks like $\mathbf{G}^{(0)}$ today. As an immediate consequence of Theorem 1, we observe that for large n, all agents' opinion dynamics will be close to their expectation. In particular, for $\mathbf{G}^{(1)}$, we have:

$$\rho_{ij} = \begin{cases} (1-p_i)(1-p_j) & \text{, if } j \in N(i) \\ 1-(1-q_i)(1-q_j) & \text{, if } j \notin N(i) \end{cases}$$

Then, we can compute the "expected" opinion dynamics for normal agents:

$$\pi_{i,t+1} = \frac{(1-p_i)(1-p_j)}{\bar{d}_i} \sum_{j \in N(i)} \pi_{j,t} + \frac{1-(1-q_i)(1-q_j)}{\bar{d}_i} \sum_{j \notin N(i)} \pi_{j,t}$$

where $\bar{d}_i = (1-p_i)\sum_{j\in N(i)}(1-p_j) + \sum_{j\notin N(i)}(1-(1-q_i)(1-q_j))$. Using standard techniques, one can compute $\lim_{t\to\infty}\pi_{i,t}$ for all agents i. Under the random evolution of the network, each agent's true limiting belief will be close to this computation for sufficiently large population sizes.

Application 2: Assortative Random Matching and Financial Volatility. Consider a financial model of contagion similar to Elliott et al. (2014) where banks hold fundamental assets of random value α_i, but also shares in each other ω_{ij}. Thus, the fundamental value of the bank can be written recursively as $\mathbf{v} = \boldsymbol{\alpha} + \boldsymbol{\Omega}\mathbf{v}$. This follows the same form of the input-output economy, so our centrality measure can be used to measure aggregate volatility. In reality, $\boldsymbol{\Omega}$ is endogenously chosen by the banks, and has been shown (e.g., Guttman (2008)) that strong banks typically invest in other strong banks, whereas weak banks typically invest in other weak banks (this is known as *assortative matching*). The question is whether this assortative matching is better for aggregate volatility relative to uniformly-at-random matching.

For large n, Theorem 1 first establishes that there *is an answer*, in the sense that with high probability one regime will *always* be more volatile than the other. Let us assume for simplicity that banks are either "strong" or "weak" and in the assortative matching regime, banks of the same type match with probability p_s and banks of different types match with probability p_d. On the other hand, in the uniformly-at-random matching regime, all banks match with the same probability p. Then we obtain values $\mathbf{v}_a, \mathbf{v}_{uar}$ for the banks with assortative matching and uniformly-at-random matching, respectively:

$$\mathbf{v}_a = (\mathbf{I} - \mathbf{P}_a)^{-1}\boldsymbol{\alpha}$$
$$\mathbf{v}_{uar} = (\mathbf{I} - \mathbf{P}_{uar})^{-1}\boldsymbol{\alpha}$$

where

$$\mathbf{P}_a = \frac{2}{n(p_s+p_d)}\begin{pmatrix} p_s\mathbf{1} & p_d\mathbf{1} \\ p_d\mathbf{1} & p_s\mathbf{1} \end{pmatrix}$$
$$\mathbf{P}_{uar} = \frac{1}{n}\mathbf{I}$$

which allows us to directly compare the volatilities of banks' fundamental values. Moreover, by Theorem 1, we can be confident that these results generalize to the case where many banks match assortatively/uniformly-at-random, even though the above analysis is done only in expectation.

Application 3: Pricing with Network Knowledge. Consider the model from before of Candogan et al. (2012), with a monopolist who chooses prices $\mathbf{p} = \{p_i\}_{i=1}^n$ to maximize profits. Each agent i has a fixed weight on how much the externality affects their own consumption; in particular, each agent has constant $\mathcal{G}_i = \sum_{j=1}^n g_{ij} < 2b_i$. Then, the network is formed randomly according to some matrix of link probabilities $\boldsymbol{\rho}$, and each agent splits its $\mathcal{G}_i$ evenly across all of its neighbors.

The question is: how much value does the realized network structure add to the monopolist's pricing decision? By Theorem 1, as $n \to \infty$, the answer is none. It suffices to only know *the network generation process*. In particular, a monopolist can set prices according to the *expected* network, and the profits from doing so arbitrarily approximate the optimal pricing with complete network knowledge.

This complements two results in the paper of Candogan et al. (2012). First, it provides a nice follow-up to Corollary 1, which states that in a symmetric network, pricing is independent of the network. Our conclusion shows that while the network of connections is symmetric, the matrix of externalities $G = \{g_{ij}\}$ is not, and therefore it might still be optimal for a monopolist to take into account network effects as the number of agents n grows large (e.g., giving discounts to central agents) when choosing his pricing vector $\mathbf{p}$. Second, this result adds to the discussion in Section 5 in Candogan et al. (2012) regarding the value of knowing the exact network structure. In special cases, where the network is drawn from a distribution of the form in Section 2 in this paper, knowledge of the realized network provides no added value over the expected network, and a monopolist can use his knowledge of the latter to extract almost all of the profits.

5 Conclusion

This paper introduces row-stochastic centrality and offers a result that simplifies the study of this centrality measure in stochastic and dynamic networks. Because the normalized adjacency matrix is asymmetric, the recent work on centrality in random networks does not apply to this centrality measure, and therefore cannot be applied to network problems that utilize this measure in their analyses.

By showing that row-stochastic centrality measures in random networks are close to their values in the average network, the stochastic and dynamic versions of several applications in the literature can be reduced to studying the average matrix. Empirical studies (e.g. Conley and Udry (2010)) routinely collect statistical information about the likelihood of link formation, and our main result shows that this information is enough to derive results that are highly robust to the realized network structure. Given the prevalence of row-stochastic centrality in many network applications, we expect that our results have broader applicability beyond the examples presented in this paper.

References

Acemoglu, Daron, Vasco Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi (2011), "The Network Origins of Aggregate Fluctuations." *Econometrica*, 80.

Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi (2015), "Systemic Risk and Stability in Financial Networks." *American Economic Review*, 105, 564–608, URL https://www.aeaweb.org/articles?id=10.1257/aer.20130456.

Ahern, Kenneth (2012), "Network Centrality and the Cross Section of Stock Returns." *SSRN Electronic Journal*.

Avella-Medina, Marco, Francesca Parise, Michael T. Schaub, and Santiago Segarra (2020), "Centrality Measures for Graphons: Accounting for Uncertainty in Networks." *IEEE Transactions on Network Science and Engineering*, 7, 520–537. Conference Name: IEEE Transactions on Network Science and Engineering.

Billingsley, Patrick (1995), *Probability and Measure*, third edition. John Wiley and Sons.

Bonacich, Phillip (1987), "Power and Centrality: A Family of Measures." *American Journal of Sociology*, 92, 1170–1182, URL https://www.journals.uchicago.edu/doi/abs/10.1086/228631. Publisher: The University of Chicago Press.

Burres, David (1985), "Homeomorphism between leontif and cobb-douglas input-output models." *Economic Letters*, 44, 49–53.

Candogan, Ozan, Kostas Bimpikis, and Asuman Ozdaglar (2012), "Optimal pricing in networks with externalities." *Operations Research*, 60, 883–905.

Charras-Garrido, Myriam and Pascal Lezaud (2013), "Extreme Value Analysis : an Introduction." *Journal de la Societe Française de Statistique*, 154, pp 66–97, URL https://hal-enac.archives-ouvertes.fr/hal-00917995.

Chung, Fan (2005), "Laplacians and the cheeger inequality for directed graphs." *Annals of Combinatorics*, 9, 1–19.

Chung, Fan and Mary Radcliffe (2011), "On the spectra of general random graphs." *Electronic Journal of Combinatorics*, 18.

Conley, Timothy G and Christopher R Udry (2010), "Learning about a new technology: Pineapple in ghana." *American Economic Review*, 100, 35–69.

Dasaratha, Krishna (2017), "Distributions of Centrality on Networks." *Games and Economic Behavior*.

Elliott, Matthew, Benjamin Golub, and Matthew O. Jackson (2014), "Financial Networks and Contagion." *American Economic Review*, 104, 3115–3153, URL https://www.aeaweb.org/articles?id=10.1257/aer.104.10.3115.

Golub, Benjamin and Matthew O Jackson (2010), "Naive learning in social networks and the wisdom of crowds." *American Economic Journal: Microeconomics*, 2, 112–49.

Guttman, Joel M (2008), "Assortative matching, adverse selection, and group lending." *Journal of Development Economics*, 87, 51–56.

Jackson, Matthew O. (2010), *Social and Economic Networks*. URL https://press.princeton.edu/books/paperback/9780691148205/social-and-economic-networks.

Mostagir, Mohamed, Asu Ozdaglar, and James Siderius (2019), "When is society susceptible to manipulation?" *Working paper*.

Parise, Francesca and Asuman Ozdaglar (2019), "Graphon games: A statistical framework for network games and interventions." *arXiv:1802.00080 [cs]*, URL http://arxiv.org/abs/1802.00080. ArXiv: 1802.00080.

Taylor, Chris (2011), "Expectation of the maximum of gaussian random variables." Mathematics Stack Exchange, URL https://math.stackexchange.com/q/89030. URL:https://math.stackexchange.com/q/89030 (version: 2011-12-06).

Yildiz, Ercan, Asuman Ozdaglar, Daron Acemoglu, Amin Saberi, and Anna Scaglione (2013), "Binary opinion dynamics with stubborn agents." *ACM Transactions on Economics and Computation (TEAC)*, 1, 19.