

RESEARCH ARTICLE

A hierarchical latent space network model for mediation

Tracy M. Sweet¹  and Samrachana Adhikari²

¹University of Maryland, College Park, MD, USA and ²New York University Grossman School of Medicine, New York, NY, USA

Corresponding author: Tracy M. Sweet, email: tsweet@umd.edu

Action Editor: Stanley Wasserman

Abstract

For interventions that affect how individuals interact, social network data may aid in understanding the mechanisms through which an intervention is effective. Social networks may even be an intermediate outcome observed prior to end of the study. In fact, social networks may also mediate the effects of the intervention on the outcome of interest, and Sweet (2019) introduced a statistical model for social networks as mediators in network-level interventions. We build on their approach and introduce a new model in which the network is a mediator using a latent space approach. We investigate our model through a simulation study and a real-world analysis of teacher advice-seeking networks.

Keywords: social network models; mediation analysis; Bayesian

1. Introduction

The goal of many large-scale studies in the social sciences is to investigate the effect of an intervention or policy on an outcome of interest, such as the effect of a new mathematics curriculum on teachers' instructional practices or the effect of vaccination advertisements on increasing vaccination rates. Such studies often involve large samples, requiring significant resources, but they are notorious for finding null results in educational studies (Coalition for Evidence-Based Policy, 2013). Further many interventions are difficult to implement or modify in the real world.

Thus, it is important that researchers implementing these large-scale studies also examine the mechanisms that underlie the intervention. Social network data may be one of the first observed intermediate outcomes before observing the primary study outcome; for example, an intervention about bullying will likely first impact how individuals interact. Second, social networks may even mediate the impacts of an intervention on individual-level outcomes; a bullying intervention may impact network structure in ways that are more conducive to improving student sense of belonging.

Social network data, however, present unique methodological challenges, since most statistical models are not appropriate for network data. This is because network ties violate the assumption of independent observations assumed by these models. Whether or not person a is friends with person b depends on what other friendships exist among the individuals in the network, whether a and b have any friends in common, and even the size of the network. To accommodate the unique interdependence present in social network data, researchers have proposed and continue to propose social network models that predict network ties (e.g. van Duijn et al., 2004; Wasserman & Pattison, 1996; Hoff et al., 2002; Snijders & Nowicki, 1997).

Although most network researchers have focused on methods for a single network, several social network models have been extended for use across multiple, independent networks; that is, multilevel extensions of these network models have been developed that can be used on network data collected across ensembles or groups of networks (e.g. Snijders & Kenny, 1999; Templin et al., 2003; Zijlstra et al., 2006; Sweet et al., 2013; Koskinen et al., 2015; Slaughter & Koehly, 2016; Lorenz et al., 2020). In addition, methods now exist to examine network-level interventions (Valente, 2012) and estimate treatment effects (e.g. Sweet et al., 2013). Moreover, Sweet (2019) proposed the idea of modeling a social network as a mediator and introduced a mixed membership model for mediation. This model assumes that an intervention impacts the subgroup structure of the networks which then impacts the outcome.

While we believe the overall idea in Sweet (2019) is promising and the proposed model has a great potential to be useful in assessing network mediation, the model is limited in that it requires networks to have clique structure. Further, the mediating variable in this model is a measure of subgroup structure, which may not always be the intermediate outcome observed as result of an intervention. Thus, alternative models are needed.

We introduce the hierarchical latent space mediation model (HLSMM), for mediation, which uses a latent space modeling approach instead of a mixed membership approach. Our model also assumes an intervention on a sample of multiple, independent networks, such as student friendship networks across schools or collaboration networks across organizations. Our model assumes that an intervention affects the network structure in treated networks, and this structure in turn affects the variability in an outcome variable in those networks. One possible application is to model the process through which group consensus is established.

Before introducing our model, we discuss modeling samples or ensembles of networks and estimating intervention effects using latent space approaches. We then introduce our model, a HLSM for mediation, and discuss its utility and estimation. We then present some evidence of feasibility through a simulation study and real-world analysis of teacher advice-seeking data.

2. Modeling ensembles of networks

Researchers have attempted to accommodate tie interdependency by either accounting for tie dependence explicitly and including network statistics as part of the model (e.g. exponential random graph models (Robins et al., 2007)) or accounting for tie dependence implicitly through the use of latent variables (Dabbs et al., 2020). We will focus on a latent variable model called a latent space model (LSM; Hoff et al., 2002), which incorporates latent space positions in the network model to account for the interdependence of network ties.

For modeling a sample of networks, we can use a HLSM (Sweet et al., 2013), a multilevel extension of LSM, although other multilevel extensions have been proposed for other network models (Snijders & Kenny, 1999; Templin et al., 2003; Zijlstra et al., 2006; Koskinen et al., 2015; Slaughter & Koehly, 2016; Lorenz et al., 2020).

A HLSM for an ensemble of binary networks denoted by adjacency matrix, A_k , for network k is given as:

$$\begin{aligned}
 A_{ijk} &\sim \text{Bernoulli}(p_{ijk}) \\
 \text{logit}(p_{ijk}) &= \beta_0 + \beta_k X_{ijk} - |Z_{ik} - Z_{jk}| \\
 \beta_k &\sim \text{MVN}_q(\mu_{\beta_k}, \Sigma_{\beta_k}) \\
 Z_{ik} &\sim \text{MVN}_d(0, \sigma_{Z_k}^2 I_d)
 \end{aligned} \tag{1}$$

where A_{ijk} denotes a tie between nodes i and j for network k , X_{ijk} is an edge-level covariate, β_0 is an intercept, and β_k is a q -dimensional vector of the associated coefficients. Parameters β_k may

vary across each network k . Without covariates, Equation (1) specifies the probability of a tie p_{ijk} between two nodes i and j in a network k as a function of the distance between the d -dimensional latent space positions Z_{ik} and Z_{jk} . The latent position Z_{ik} is assumed to have multivariate normal distribution with mean 0 and diagonal variance–covariance matrix $(\sigma_{Z_k}^2 I_d)$, where I_d is a d -dimensional identity matrix.

Networks with larger values of the variance of the latent position, $\sigma_{Z_k}^2$, would generally have latent positions that are more dispersed than networks with smaller values of $\sigma_{Z_k}^2$. Because node dispersion is generally related to increases in distance between pairs of nodes, $\sigma_{Z_k}^2$ is correlated with the density of the network. Given a collection of networks, it may be of interest to explore or even model the differences in the variance of the latent space positions across networks.

2.1 Modeling interventions with LSMs

Consider an intervention in which entire networks are assigned to treatment or condition. Because we will eventually model the effect of the intervention as a function of σ_Z^2 in Equation (1), we first demonstrate how σ_Z^2 impacts network structure.

Recall that in a LSM, each network is represented by relative positions of its nodes in a low-dimensional space and the probability of a tie is a function of the distance between the nodes. Thus, σ_Z^2 controls the spread or scale of the distance, and networks with smaller values of σ_Z^2 will have higher density than networks with larger values of σ_Z^2 , because smaller values of σ_Z^2 map to smaller pairwise distances and higher tie probabilities. In addition, we also expect to see an increase in the number of isolated nodes as σ_Z^2 increases.

However, unlike random graph models, LSMs are also known to preserve reciprocity and transitivity (Hoff et al., 2002). This is because, as long as the position of a node relative to the neighboring nodes is preserved, reciprocity and transitivity will also be preserved. Therefore, with the exception of extreme differences in σ_Z^2 , reciprocity and transitivity should be similar across values of σ_Z^2 .

For illustration, we simulated 100 networks for values of σ_Z^2 and plotted the median network statistic value (orange circle) along with 0.025 and 0.975 quantiles, indicated by the line segment. Figure 1 illustrates how the variance of the latent space positions in a LSM impacts density, variance in degree, the number of isolated nodes, and mean geodesic distance (among connected nodes). The plot showing geodesic distance is misleading. The distances are not actually decreasing with increasing σ_Z^2 ; rather, the number of isolated nodes who are not connected to other nodes is increasing, and these distances are not included.

Aside for extreme differences in σ_Z^2 , there is little impact on the levels of reciprocity and transitivity although we note that reciprocity and transitivity naturally increase as density increases. Here, we use σ_Z^2 values 2 to 10 to illustrate a wide range, but note that in Section 4.1, we consider networks generated from LSMs with expected values of σ_Z^2 of 2.7 in the control condition.

Thus, an intervention that impacts σ_Z^2 would change the overall density and the number of isolated individuals in the network, while preserving some standard network structures such as transitivity and reciprocity. For example, in an intervention focused on improving school climate, we may expect greater unity and consensus among its teachers, but we would not expect clique structure or reciprocity to change. Similarly, teachers who were isolated previously may be now be connected to someone in the network as a result of a decrease in σ_Z^2 .

3. A hierarchical latent space mediation model

Let us consider an example that will motivate our proposed model. Suppose, there is a cluster-randomized trial in which schools are randomly assigned to an intervention and that the

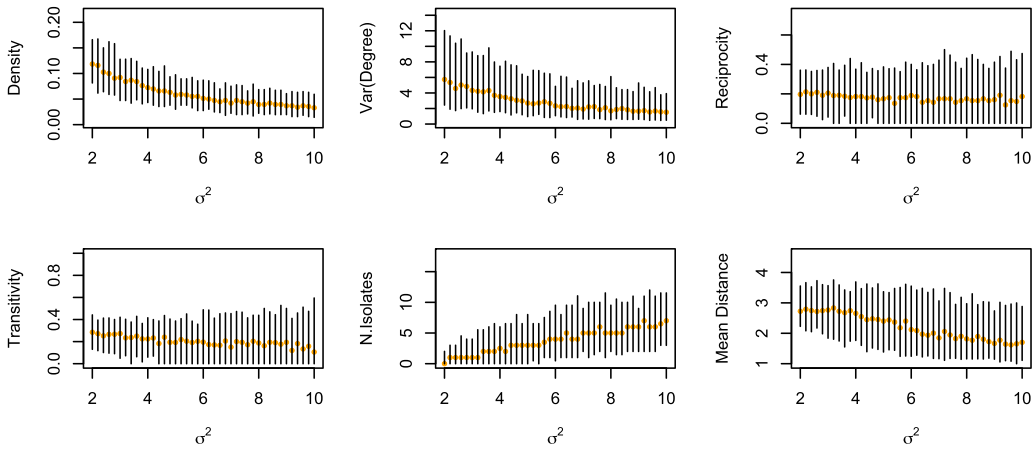


Figure 1. Network statistics from networks generated from latent space models where σ_z^2 ranges from 2 to 10 to illustrate the impact of σ_z^2 on density and the number of isolated nodes with comparatively little impact on reciprocity and transitivity.

intervention changes how teachers within each school interact. Let us also assume the goal of this intervention is to improve teacher buy-in about an instructional method. We hypothesize that treated schools have different types of interactions which result in a stronger consensus about that instructional method.

To fit a mediation model, a classic approach is to use the three regression equations shown in Equation (2) and made popular by Baron & Kenny (1986). We present the Bayesian version of these equations where generally the last two lines are estimated, and the mediation or indirect effect is defined as $\beta_2\gamma$ (Yuan & MacKinnon, 2009). The model is given as

$$\begin{aligned}
 Y_k &\sim N(\alpha_1 + \beta_1 T_k, \tau_1) \\
 M_k &\sim N(\alpha_2 + \beta_2 T_k, \tau_2) \\
 Y_k &\sim N(\alpha_3 + \beta_3 T_k + \gamma M_k, \tau_3)
 \end{aligned}
 \tag{2}$$

where k indexes the observed unit of randomization, Y_k denotes the outcome, T_k denotes the treatment status, and M_k denotes the mediator. Similarly, $\alpha_1, \alpha_2,$ and α_3 are intercept parameters, $\beta_1, \beta_2,$ and β_3 are coefficients associated with treatment indicator T_k , and γ is the coefficient associated with mediator M_k .

In our context, our observed units are the networks, and our outcome of interest is teacher consensus, measured as the variability in opinions about this instructional method. Further, we assume that the intervention is focused on bringing teachers together and increasing collaboration. We present this scenario as a Bayesian mediation model where we replace M_k with a parameter from a HLSM for network k and incorporate the last two lines of Equation (2).

Thus, the hierarchical latent space mediation model (HLSMM) model is given as:

$$\begin{aligned}
 A_{ijk} &\sim \text{Bernoulli}(p_{ijk}) \\
 \text{logit}(p_{ijk}) &= \beta_0 - |Z_{ik} - Z_{jk}| \\
 Z_{ik} &\sim \text{MVN}_d(0, \sigma_{Z_k}^2 I_d) \\
 \log(\sigma_{Z_k}) &\sim N(\gamma_0 + \gamma_1 T_k, \tau_Z) \\
 Y_{ik} &\sim N(\mu, \sigma_{Y_k}^2) \\
 \log(\sigma_{Y_k}) &\sim N(\theta_0 + \theta_1 \log(\sigma_{Z_k}) + \nu T_k, \tau_Y)
 \end{aligned}
 \tag{3}$$




Figure 2. The directed acyclic graph for the HSLM for mediation illustrates the relations among the independent variable T , the mediating parameter $\log(\sigma_Z)$, and the outcome variable $\log(\sigma_Y)$.

where the treatment indicator is still T_k ; i and j index the individuals within each network and k indexes the network. For any network, A_k is generated from a LSM such that the variance-covariance matrix of the latent space positions is $\sigma_{Z_k}^2 I_d$. Note that pretreatment confounders can be included in the model as well.

We can also represent this model as the directed acyclic graph given in Figure 2; the mediation analysis is illustrated among treatment T , the mediator $\log(\sigma_Z)$, and the outcome variable $\log(\sigma_Y)$. Compared to Equation (2), the mediation analysis is also tied to network model such that a network tie from node i to node j in network k (A_{ijk}) is a function of latent space positions Z_{ik} and Z_{jk} with variance $\sigma_{Z_k}^2$.

Note that Figure 2 and Equation (3) denote the impact of an independent variable and mediating variable on the variability of some outcome $\log(\sigma_Y)$, but other outcome variables could be used. Because the mediator is the log standard deviation of the latent space positions, a natural outcome is some measure of consensus or the log standard deviation in the outcome variable among individuals in each network. Also note that Equation (3) is parameterized such that the outcome and mediator both have normal distributions, similar to Equation (2), but we could parameterize this model so that σ_Z and σ_Y have lognormal distributions.

Returning to our contextual example, if the intervention is effective, the treated networks will have a smaller latent position variance. Decreasing the log standard deviation of the latent space positions, $\log(\sigma_Z)$, essentially decreases the number of isolated nodes and increases the overall density of the network. Not having any isolated nodes and increasing the number of ties would therefore result in higher levels of connection in the network. Thus, we hypothesize that well-connected networks have higher levels of consensus than networks with sparse connections.

3.1 Estimation

The HLSMMs in this manuscript were fit using a Markov chain Monte Carlo (MCMC) algorithm (Gelman et al., 2013), and we used a combination of Metropolis-Hastings and Metropolis updates within Gibbs sampling. The Metropolis/Gibbs algorithm is standard and allows any reader familiar with MCMC to understand the estimation process as well as augment the code to expand or adapt the proposed model for their purposes. However, the HLSMM can be fit using a variety of MCMC algorithms as well as posterior approximation methods such as variational inference.

Moreover, the MCMC algorithm combines the network portion of the model and the mediation model and updates all parameters simultaneously. That is, the latent positions given the

observed networks are updated iteratively alongside the regression coefficients of the mediation portion of the model. Fitting the HLSMM in this way accounts for the estimation error in the latent positions when estimating the regression coefficients used to calculate the indirect effect.

Given the normal distributions specified in the model for $\log(\sigma_Z)$ and $\log(\sigma_Y)$, we used Gibbs updates for γ_0 , γ_1 , θ_0 , θ_1 , and ν , using normal distributions as conjugate priors. Similarly, for variance hyper-parameters τ_Z and τ_Y , we employed conjugate inverse-Gamma distributions so that they were also updated via Gibbs. The latent space positions as well as the variance of the latent space positions (reparametrized as $\log(\sigma_Z)$) were each updated using a random walk Metropolis-Hastings update. We used conjugate priors when possible, but other prior distributions can be used. For example, half-Cauchy distributions could replace inverse-Gamma distributions for the variance updates.

Note that LSMs have two identifiability issues. The first is that there are an infinite number of latent space position configurations possible due to translation, reflection, and rotation. A procrustes transformation as used in Hoff et al. (2002) can be applied to obtain estimates of the latent space positions with a stable configuration. Note that this transformation can be applied post-sampling as convergence can be obtained without it.

The second identifiability issue is discussed less often but has larger impacts on the HLSMM parameters. There is an identifiability between β_0 , which can be thought of as an intercept and related to the overall probability of a tie, and the variance of the latent space positions $\sigma_{Z_k}^2$, which also impacts tie probability through the dispersion of nodes in the latent space. We recommend fixing β_0 when estimating $\sigma_{Z_k}^2$ is of particular importance.

4. Empirical examples

As a proof of concept, we explore some of the operating characteristics and demonstrate the feasibility of the HLSMM with two empirical examples. The first is a series of simulations to help readers connect mediation parameter values to variance parameters and network structures. Through these simulations, we also examine the conditions under which the proposed model can be most useful. The second example is a real-world application of the HLSMM that involves advice-seeking networks among school staff at 14 schools in 2010 and 2015.

4.1 Simulation study

The purpose of this small simulation study is to examine the influence of various parameters on both parameter recovery and model utility. We consider nine different combinations of parameters γ_1 and θ_1 in Equation (3). Recall that γ_1 is the effect of the intervention T on the log standard deviation of the latent space positions $\log(\sigma_Z)$, so we explore settings where the intervention has three different effects on this measure of variability. Similarly, θ_1 is the conditional effect of $\log(\sigma_Z)$ on $\log(\sigma_Y)$ given T , and we consider three different values for this parameter as well.

We selected values of γ_1 that resulted in differences in $\log(\sigma_Z)$ that produced visually distinct but still realistic networks and values of θ_1 that yielded analogous values for $\log(\sigma_Y)$. Note that negative values of γ_1 result in smaller variance of latent space positions. Our positive value of θ_1 maps to a positive relationship between latent space variability $\log(\sigma_Z)$ and variability in the outcome $\log(\sigma_Y)$.

Table 1 summarizes the nine cells of our simulation, and each cell represents one combination of values for γ_1 and θ_1 . We also include sample variances (across replications) to illustrate the indirect and direct effects on the outcome; note that σ_Z^2 and σ_Y^2 are given instead of $\log(\sigma_Y)$ and $\log(\sigma_Y)$ for ease of interpretation. The full data generating model used is given as:

Table 1. A summary of parameter values across the nine cells of the simulation study along with approximate sample variances

Setting/parameters	γ_1	θ_1	Avg. $\sigma_{Y_k}^2$, control	Avg. $\sigma_{Y_k}^2$, treated
1	-0.2	0.5	0.67	0.45
2	-0.2	0.7	0.82	0.51
3	-0.2	0.9	1	0.57
4	-0.4	0.5	0.67	0.37
5	-0.4	0.7	0.82	0.38
6	-0.4	0.9	1	0.40
7	-0.6	0.5	0.67	0.30
8	-0.6	0.7	0.82	0.29
9	-0.6	0.9	1	0.28

$$\begin{aligned} \log(\sigma_{Z_k}) &\sim N(0.5 + \gamma_1 T_k, 0.09), k = 1, \dots, 30 \\ Z_{ik} &\sim MVN_2(0, \sigma_{Z_k}^2 I_2), i = 1, \dots, 20 \\ p_{ij} &\sim \frac{\exp(-|Z_{ik} - Z_{jk}|)}{1 + \exp(-|Z_{ik} - Z_{jk}|)}, i, j = 1, \dots, 20 \\ A_{ijk} &\sim \text{Bernoulli}(p_{ij}), i, j = 1, \dots, 20 \\ \log(\sigma_{Y_k}) &\sim N(-0.45 + \theta_1 \log(\sigma_{Z_k}) - 0.1T_k, 0.09) \\ Y_{ik} &\sim N(0, \sigma_{Y_k}^2) \end{aligned} \tag{4}$$

where k indexes the 30 networks each with 20 nodes. The independent variable T_k indicates whether network k is in the treated group such that 15 networks are treated and 15 networks are not.

In Equation (4), the expected $\log(\sigma_Z)$ for control networks and treated networks is defined as γ_0 and $\gamma_0 + \gamma_1$, respectively. Thus, in selecting values for each, our goal was to generate networks that reflected real-world data. The expected σ_Z^2 under the control condition is approximately 2.72 and under the three treated conditions ($\gamma_1 = -0.2, -0.4, -0.6$) is 1.82, 1.22, and 0.82 respectively.

Figure 3 shows 30 networks generated from Equation (4) where $\gamma_0 = 0.5$ and $\gamma_1 = -0.4$. The control networks are shown on the top three rows of plots and the treated networks on the bottom three rows. Overall, the control networks tend to have more isolated nodes and are less dense than the treated networks. Thus, $\gamma_1 = -0.2$ would result in less noticeable differences between the conditions, and $\gamma_1 = -0.6$ would result in more noticeable differences between conditions.

We selected values for θ_1 in a similar manner aiming for σ_Y to be distinct enough to detect differences. Again, we selected error variances ($\tau = 0.09$) to observe adequate separation of distributions of $\log(\sigma_{Y_k})$ to detect effects. Finally, we also chose to include a nonzero value for $\nu = -0.1$ to indicate a partial mediation. Figure 4 shows a plot of $\log(\sigma_{Y_k})$ against $\log(\sigma_{Z_k})$ for the 30 networks generated from Equation (4) where $\gamma_0 = 0.5, \gamma_1 = -0.4, \theta_0 = -0.45, \theta_1 = 0.9$, and $\nu = 0.1$. There is evidence of a positive effect of $\log(\sigma_Z)$ on $\log(\sigma_Y)$.

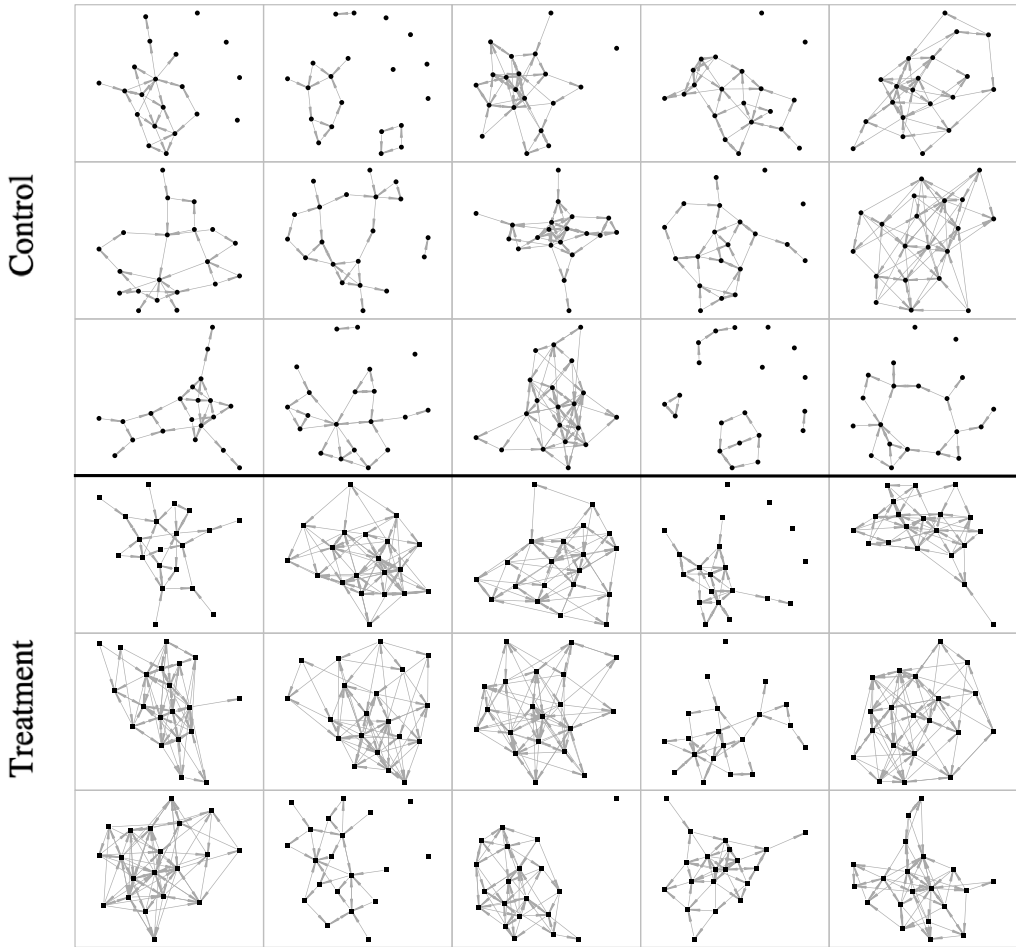


Figure 3. Networks simulated from a HLSM for Mediation where T indicated treatment and the treated networks have smaller σ_{Z_k} (see Equation (4)).




Figure 4. The scatterplot of $\log(\sigma_\gamma)$ against $\log(\sigma_z)$ shows a positive effect of the mediator on the outcome.

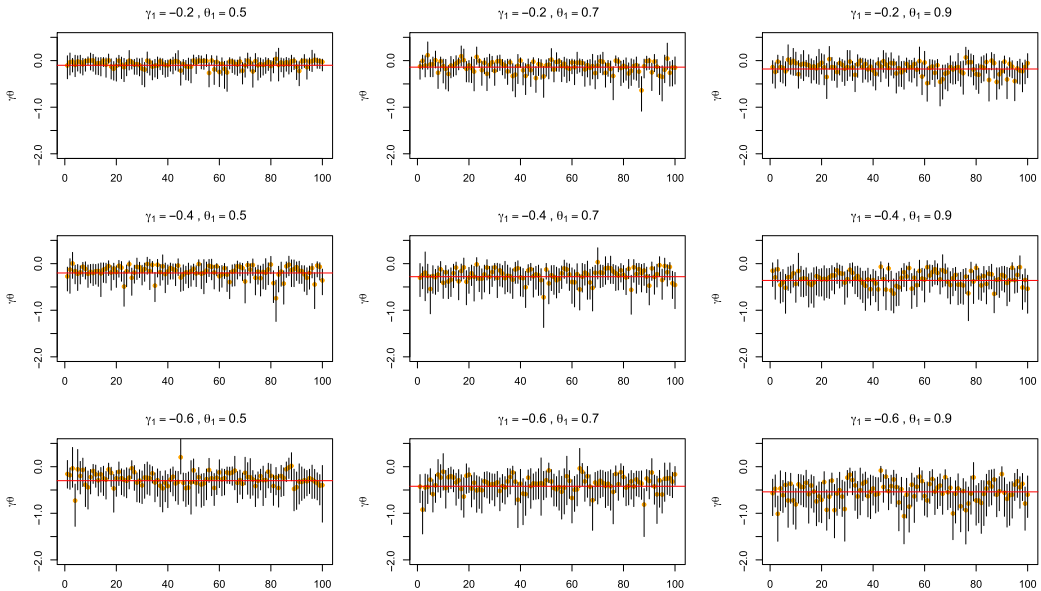


Figure 5. Simulation summary of mediation effect corresponding to settings 1–9 in Table 1. The lower and upper bound of each error bar indicates the 95% highest posterior density credible interval for the posterior distribution of the average causal mediation effect and the circle represents MAP of the distribution. The red horizontal line represents the true value of the parameters used to generate the data.

Within each simulation condition given in Table 1, we simulated data from a HLSMM generating model, Equation (4), and fit a HLSMM, Equation (3), using the following priors:

$$\begin{aligned}
 \gamma_0, \theta_0 &\sim N(0, 100) \\
 \gamma_1, \theta_1, \nu &\sim N(0, 9) \\
 \tau_Z, \tau_Y &\sim IG(5, 0.5)
 \end{aligned}
 \tag{5}$$

We selected weakly informative priors for γ_0 and θ_0 and chose more informative priors for $\gamma_1, \theta_1,$ and ν and error variances τ_Z and τ_W . (See the Appendix for a prior sensitivity analysis.) For each model fit, we ran three chains of length 10,000; we removed the first 1,000 draws and thinned each chain by 35 to obtain a final posterior sample of 774; this length and thinning provided adequate convergence and autocorrelation across all parameters. This process was repeated 100 times for each cell of the simulation (Table 1).

To evaluate model performance under each setting, we report the coverage rate, defined as the proportion of replications in which the 95% equal-tailed credible interval covered the data generating parameter. In addition, we also report the posterior probabilities that the mediation effect $\gamma_1\theta_1 < 0$. Note that the mediation effect is negative for all cells of the simulation.

Figure 5 summarizes the posterior distributions of the mediation effect (indirect effect) $\gamma_1\theta_1$ of the 100 replications across all 9 conditions. The circles represent the maximum a posteriori (MAP) estimates, and the bars indicate the 95% highest posterior density (HPD) credible intervals. The coverage rates for the mediation effect are between 0.91 and 0.98, indicating that the indirect effect coverage is as expected. In addition, parameter coverage for all parameters estimated in the model range from 0.91 to 1.00. Furthermore, we note that the posterior variance for $\gamma_1\theta_1$ increases as the effect sizes of γ_1 and θ_1 increase.

Table 2. Mean posterior probability that $\gamma_1\theta_1 < 0$ across all nine simulation conditions

γ_1	θ_1	$P(\gamma_1\theta_1 < 0)$
-0.2	0.5	0.81
-0.2	0.7	0.86
-0.2	0.9	0.88
-0.4	0.7	0.94
-0.4	0.5	0.96
-0.4	0.9	0.98
-0.6	0.5	0.94
-0.6	0.7	0.98
-0.6	0.9	0.99

We also examine the posterior probability of the indirect effect being less than 0. We find that as the true indirect effect becomes more negative (a stronger effect), the posterior probability increases. It appears that changing the effect sizes of γ_1 has a larger impact than θ_1 on these posterior probabilities. The exact probabilities are given in Table 2.

4.2 Application to teacher networks

We now present an application of the HLSMM to real-world data. These data come from a longitudinal study of workplace and social interactions among school staff that are part of the Distributed Leadership Studies at Northwestern University (distributedleadership.org). These data represent language arts advice-seeking relationships among school staff in 14 elementary schools in 2010 and 2015 in a suburban school district in the midwestern part of the US. There were major policy changes during this time including a new mathematics curriculum, the use of mathematics instructional coaches for teachers, and high stakes testing in mathematics.

While advice-seeking ties around mathematics have been studied in detail (Spillane et al., 2017, 2015, 2018), language arts advice-seeking ties have not. We aim to examine the effects of the policy changes in mathematics on the language arts advice-seeking networks, and whether these networks then mediated consensus about teachers' influence on school policy.

School staff in these 14 schools were asked "During this school year, to whom have you turned for advice and/or information about curriculum, teaching, and student learning?" For each person nominated, respondents then classified the connection as being about language arts, math, or science. The language arts networks for the years 2010 and 2015 are shown in Figures 6 and 7, respectively. Note that school staff includes teachers and administrators as well as instructional support teachers (math coaches) but does not include all staff employed by the school district such as nurses, administrative assistants, custodians, or para-educators.

Response rates varied by school but on average were approximately 85%. Respondents who did not submit a survey were excluded from the network unless they were nominated by at least two other staff members, and these individuals were assumed to have no outgoing ties. Further, staff members who took the survey but did not nominate anyone remained in the networks as having no outgoing ties.

As part of this study, school staff were asked a number of items about mathematics and mathematics pedagogy, but there were a few survey items that were not subject-specific. One is a series of items asking respondents "How much influence do teachers have over school policy in each

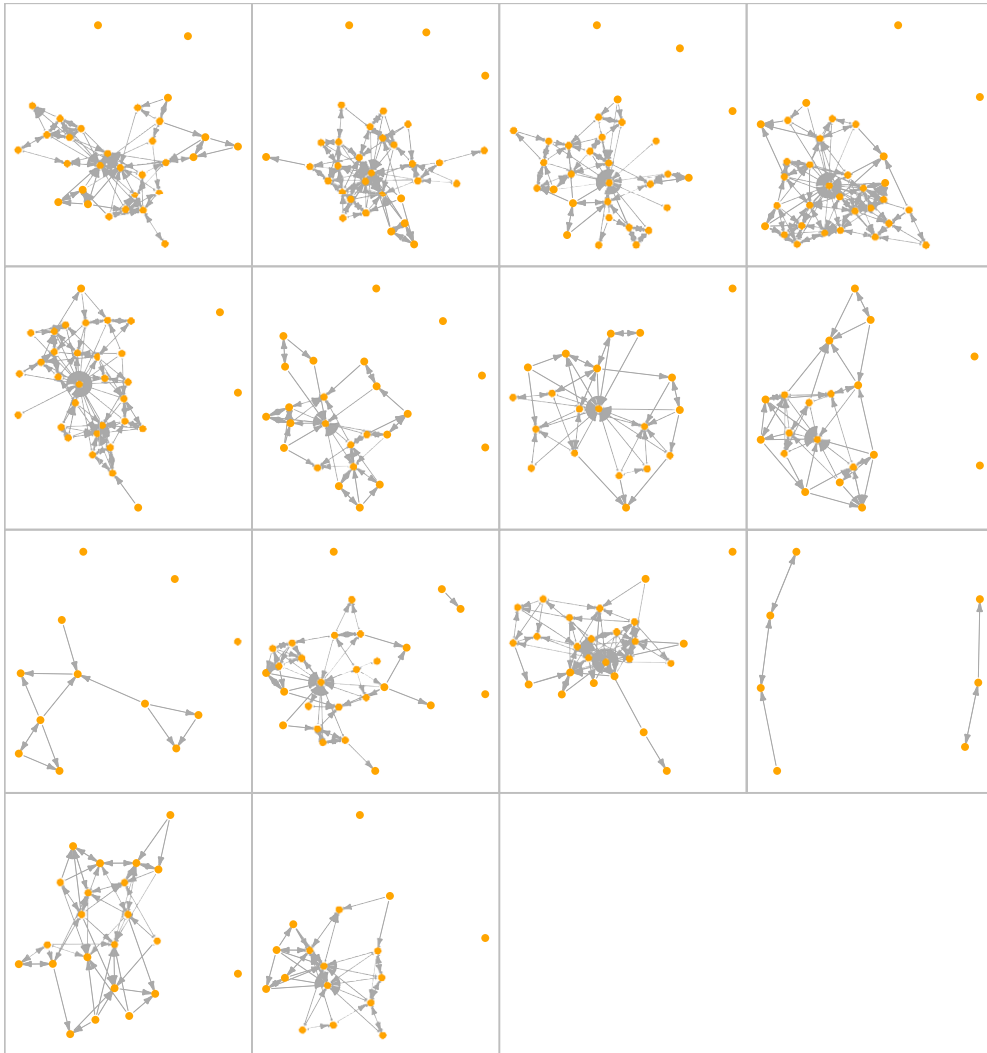


Figure 6. Language arts advice-seeking networks among school staff in 14 elementary schools in 2010; note that response rates in one school were particularly low.

of the areas below?” and the areas listed were as follows: hiring professional staff; planning how discretionary funds should be used; determining which books and instructional materials are used in classrooms; establishing the curriculum and instruction program; determining the content of in-service programs; setting standards for student behavior; and determining goals for improving the school.

To demonstrate how the HLSMM could be used in practice, we examine the extent to which language arts advice-seeking networks mediated the effect of the district changes on consensus about teacher’s influence on school policy. The belief is that structural differences in the language arts networks in 2010 compared to 2015 may impact differences in school staff consensus about teacher influence over policy (in 2010 compared to 2015).




Figure 7. Language arts advice-seeking networks among school staff in 14 elementary schools in 2015.

The HLSMM fit to these data is given as:

$$\begin{aligned}
 A_{ijk} &\sim \text{Bernoulli}(p_{ij}), i, j = 1, \dots, n_k \\
 \text{logit}(p_{ij}) &\sim \beta_0 - |Z_{ik} - Z_{jk}| \\
 Z_{ik} &\sim \text{MVN}_2(0, \sigma_{Z_k}^2 I_2), i = 1, \dots, n_k \\
 \log(\sigma_{Z_k}) &\sim N(\gamma_0 + \gamma_1 T_k, \tau_Z) \\
 \log(\sigma_{Y_k}) &\sim N(\theta_0 + \theta_1 \log(\sigma_{Z_k}) + \nu T_k, \tau_Y) \\
 \gamma_0, \theta_0 &\sim N(0, 100) \\
 \gamma_1, \theta_1, \nu &\sim N(0, 9) \\
 \tau_Z, \tau_Y &\sim \text{IG}(5, 0.5)
 \end{aligned} \tag{6}$$




Figure 8. Density plots of the posterior distributions from HLSMM fit to teacher advice-seeking data. The distribution of the mediation effect $\gamma_1\theta_1$ suggests very little evidence that the network structure mediates changes in teacher consensus.

where $A_{ijk} = 1$ if teacher i seeks language arts advice from teacher j in school k and 0 otherwise, $T_k = 1$ if the network is from 2015 and 0 otherwise, and σ_{Y_k} is the standard deviation in teacher influence scores for school k .

Because our model hinges on accurate estimation of σ_{Z_k} , we chose to fix β_0 . To fit the HLSMM, we ran a single chain of length 200,000. We broke the chain into three pieces to assess convergence; the potential scale reduction factor for the mediation effect was 1 with an upper confidence bound of 1.01, indicating adequate mixing. After tossing the first 1000 steps and thinning by every 175th step, our posterior sample size was 1138.

The posterior distributions for γ_1 , θ_1 , ν , and $\gamma_1\theta_1$ are shown in Figure 8, where the HPD credible intervals are indicated by the shaded region. There is evidence of an intervention effect γ_1 on the log (σ_Z), and the posterior probability that $\gamma_1 > 0$ is 0.94. The evidence of a mediated or indirect effect of log (σ_Z) is weak; the posterior probability that $\gamma_1\theta_1 < 0$ is 0.71, which is not statistically significant from a frequentist perspective that offers little evidence of a mediated effect. Further, the posterior probability that $\nu < 0$ is 0.93 which suggests some evidence of a direct effect.

Thus, we conclude that there is evidence that the networks in 2015 compared to those in 2010 have a larger latent space variance in a LSM. There is some evidence of a direct effect of policy changes on staff consensus about their influence, but little evidence that the spread of the latent space positions mediated the effect of policy changes on teacher consensus.

This example serves as an illustration of our model as opposed to a formal analysis of these networks. Our evidence was quite weak, but even if our finding for a mediated effect could be considered statistically significant, there were too many limitations for these results to be used in practice. We are comparing networks in 2010 with networks in 2015, and there may be other causes for differences in network structure over the years than just the district-led policy changes. Similarly, there may be other variables that impact differences in teacher consensus.

As an illustration, however, we can continue with this real-world data analysis to examine goodness of fit. Note that we did fix β_0 as there is potentially an identifiability issue between β_0 and σ_{Z_k} . We selected several possible values for $\beta_0 = (-3, -2, -1, 0, 1, 2, 3)$. If we fixed the intercept at a value that was too small, the overall probability of tie would be not be well recovered because the model is defined by subtracting distance between the latent positions from the intercept. Subtracting a number from a small number would yield results in fitted model that predicts extremely few ties. Similarly, if we selected an intercept that was too large, the latent positions would need to be quite far and it is possible that our model would produce networks dissimilar to our observed data. We hypothesized that $\beta_0 = 1$ would be ideal and decided to rely on goodness-of-fit checks to confirm.

We use posterior predictive checks and generate network data from the parameters specified at each step in our MCMC chain. We examine goodness of fit in two stages: we first investigate the network portion of the model and then we examine the regression portion of the model. For the network model fit, we use the latent space positions estimated at each step of the chain to generate




Figure 9. Posterior predictive checks for the latent space portion of the HLSMM show that network density, reciprocity, and transitivity are well recovered.

28 networks and evaluate our model fit using density, reciprocity, and transitivity in each of the 28 networks.

Figure 9 shows the distributions of density, reciprocity, and transitivity for each network generated from parameters at each step in the MCMC chain. The observed network statistic for each network are shown with a vertical line. In general, density, reciprocity, and transitivity are all well recovered, indicating that the LSM was able to capture these aspects of the observed network.

Note that posterior predictive checks using values of $\beta_0 = -1, -2$ resulted in models that generated networks with lower densities, reciprocity, and transitivity values than in our observed

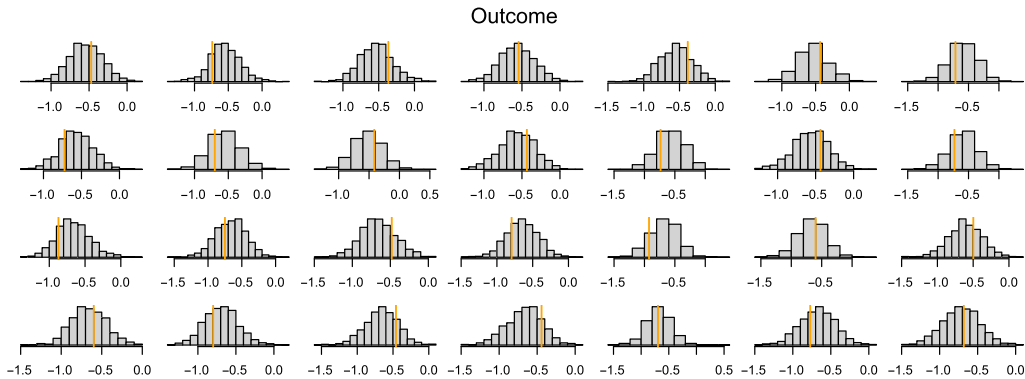


Figure 10. Posterior predictive checks for the regression portion of the HLSMM show that the predicted outcome $\log(\sigma_Y)$ in each network is well recovered.

network; similarly values of $\beta_0 = 3$ produced networks that tended to generate networks higher in reciprocity and transitivity than in our observed data. Using $\beta_0 = 0$ or $\beta_0 = 2$ did not produce extremely different networks, although reciprocity recovery was slightly lower for $\beta_0 = 0$.

For the regression portion of the model fit, we took the parameters at each step of the MCMC chain to generate outcome data $\log \sigma_{Y_k}$. We show the distribution of our predicted outcome data with the observed data in Figure 10. The posterior predictive checks suggest the outcome data were well recovered.

Thus, we find the HLSMM fitted to the language arts advice-seeking network data to be adequate because the HLSMM generates data quite similar to our observed data. Despite the lack of evidence of mediation, we believe this example illustrates the feasibility of our model and demonstrates how posterior predictive checks can be used to assess model fit.

5. Discussion

To contribute to the methodological research exploring how social networks could act as mediators in interventions on networks, we introduced the HLSMM that embeds a LSM into a mediation analysis.

Our model is specifically aimed at interventions that increase or decrease connectedness (average centrality and density) and whose network structure is believed to impact a network-level outcome. The model we presented uses the variance of the latent positions in a network model to mediate the effect of a network-level treatment on the variance of an outcome.

We also illustrated the feasibility of our model through a small simulation study showing that our model recovers parameters and can estimate mediation effects. We also presented a real-world example of a natural experiment on networks. Although we did not find strong evidence of a mediation effect, this example showed how this model can be applied to data and how posterior predictive checks can be used for assessing goodness of fit.

We note that many mediation analyses are aimed at estimating a causal effect. We argue that the HLSMM could be used in a causal mediation analysis in the same way that the mediation model put forth by Baron & Kenny (1986) can be used in a causal mediation analysis. Causal inference requires a series of assumptions, and our model's ability to meet these assumptions is context- and data-dependent.

Further, this model does not account or accommodate within-network interactions, otherwise known as network interference or peer effects. How individuals within each network are

influencing each other and their outcomes is not being modeled; this is a limitation of this model and other multilevel network models. Rather this model is looking at a mediation effect on a network-level outcome.

Thus, a natural next step in this work is to examine how network-level interventions and mediating variables could impact node-level outcomes. Being able to map changes across and within networks will help to improve our understanding of how individual outcomes are shaped through interactions. We imagine extensions to this model would then account for network interference/peer influence and align with current research in that area (e.g. An, 2018; VanderWeele & An, 2013; Ogburn, VanderWeele et al., 2017; Manski, 2000).

Acknowledgments. We thank Jim Spillane and the Distributed Leadership Studies at Northwestern University (<http://www.distributedleadership.org>), including the NebraskaMATH study, for the data used in this paper.

Competing interests. None.

References

- An, W. (2018). Causal inference with networked treatment diffusion. *Sociological Methodology*, 48, 152–181.
- Baron, R. M., & Kenny, D. A. (1986). The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology*, 51, 1173.
- Coalition for Evidence-Based Policy (2013). Randomized Controlled Trials Commissioned by the Institute of Education Sciences Since 2002: How Many Found Positive Versus Weak or No Effects, Tech. rep., Coalition for Evidence-Based Policy.
- Dabbs, B., Adhikari, S., & Sweet, T. (2020). Conditionally Independent Dyads (CID) network models: A latent variable approach to statistical social network analysis. *Social Networks*, 63, 122–133.
- Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (2013). *Bayesian data analysis* (3rd ed.). Boca Raton, FL: Taylor & Francis.
- Hoff, P. D., Raftery, A. E., & Handcock, M. S. (2002). Latent space approaches to social network analysis. *Journal of the American Statistical Association*, 97, 1090–1098.
- Koskinen, J., Caimo, A., & Lomi, A. (2015). Simultaneous modeling of initial conditions and time heterogeneity in dynamic networks: An application to foreign direct investments. *Network Science*, 3, 58–77.
- Lorenz, G., Boda, Z., Salikutluk, Z., & Jansen, M. (2020). Social influence or selection? Peer effects on the development of adolescents' educational expectations in Germany. *British Journal of Sociology of Education*, 41, 643–669.
- Manski, C. F. (2000). Economic analysis of social interactions. *Journal of Economic Perspectives*, 14, 115–136.
- Ogburn, E. L., VanderWeele, T. J., et al. (2017). Vaccines, contagion, and social networks. *Annals of Applied Statistics*, 11, 919–948.
- Robins, G., Pattison, P., Kalish, Y., & Lusher, D. (2007). An introduction to exponential random graph (p^*) models for social networks. *Social Networks*, 29, 173–191, special Section: Advances in Exponential Random Graph (p^*) Models.
- Slaughter, A. J., & Koehly, L. M. (2016). Multilevel models for social networks: Hierarchical Bayesian approaches to exponential random graph modeling. *Social Networks*, 44, 334–345.
- Snijders, T., & Kenny, D. (1999). The social relations model for family data: A multilevel approach. *Personal Relationships*, 6, 471–486.
- Snijders, T., & Nowicki, K. (1997). Estimation and prediction for stochastic blockmodels for graphs with latent block structure. *Journal of Classification*, 14, 75–100.
- Spillane, J., Hopkins, M., & Sweet, T. (2015). Intra- and inter-school instructional interactions: Exploring conditions for instructional knowledge production within and between schools. *American Journal of Education*, 122, 71–110.
- Spillane, J. P., Hopkins, M., & Sweet, T. M. (2018). School district educational infrastructure and change at scale: Teacher peer interactions and their beliefs about mathematics instruction. *American Educational Research Journal*, 55, 532–571. doi: 10.3102/0002831217743928.
- Spillane, J. P., Shirrell, M., & Sweet, T. M. (2017). The elephant in the schoolhouse: The role of propinquity in school staff interactions about teaching. *Sociology of Education*. doi: 10.1177/0038040717696151.
- Sweet, T. M. (2019). Modeling social networks as mediators: A mixed membership stochastic blockmodel for mediation. *Journal of Educational and Behavioral Statistics*, 44, 210–240.
- Sweet, T. M., Thomas, A. C., & Junker, B. W. (2013). Hierarchical network models for education research: Hierarchical latent space models. *Journal of Educational and Behavioral Statistics*, 38, 295–318.
- Templin, J., Ho, M.-H., Anderson, C., & Wasserman, S. (2003). Mixed effects p^* model for multiple social networks. In *Proceedings of the American Statistical Association: Bayesian Statistical Sciences Section* (pp. 4198–4024). Alexandria, VA: American Statistical Association.
- Valente, T. W. (2012). Network interventions. *Science*, 337, 49–53.

van Duijn, M., Snijders, T., & Zijlstra, B. (2004). p2: A random effects model with covariates for directed graphs. *Statistica Neerlandica*, 58, 234–254.

VanderWeele, T. J., & An, W. (2013). Social networks and causal inference. In S. Morgan (Ed.), *Handbook of Causal Analysis for Social Research* (pp. 353–374). New York, NY: Springer.

Wasserman, S., & Pattison, P. (1996). Logit models and logistic regressions for social networks: I. An introduction to Markov graphs and p*. *Psychometrika*, 61, 401–425.

Yuan, Y., & MacKinnon, D. P. (2009). Bayesian mediation analysis. *Psychological Methods*, 14, 301.

Zijlstra, B., van Duijn, M., & Snijders, T. (2006). The multilevel p2 model. *Methodology: European Journal of Research Methods for the Behavioral and Social Sciences*, 2, 42–47.

Appendix

Prior sensitivity analysis

We conducted a prior sensitivity analysis because prior distribution specification may be challenging with Bayesian models. We considered four different prior distributions as detailed in Table 3 where the first row is considered to be standard. We then vary the priors for the other parameters to be more or less informative.

For each replication of our sensitivity analysis, we changed the prior distribution for a single parameter and kept the priors for all other parameters to be those given in Table 3. For example, when exploring the priors for γ_0 , the priors for γ_1 , θ_0 , θ_1 , ν , τ_Z , τ_Y did not change and remained fixed to the distribution given in the first row of Table 3. Thus, our sensitivity analysis had 28 cells, and for each cell in the table, we ran 10 replications.

The data generating model is given as:

$$\begin{aligned} \log(\sigma_{Z_k}) &\sim N(0.5 - 0.6X_k, 0.09) \\ \log(\sigma_{Y_k}) &\sim N(-0.45 + 0.5 \log(\sigma_{Z_k}) - 0.1X_k, 0.09) \\ A_{ijk} &\sim \text{Bernoulli}(p_{ij}) \\ p_{ij} &\sim \frac{\exp(-|Z_{ik} - Z_{jk}|)}{1 + \exp(-|Z_{ik} - Z_{jk}|)} \\ Z_{ik} &\sim \text{MVN}_d(0, \sigma_{Z_k}^2 I_2) \\ Y_{ik} &\sim N(0, \sigma_{Y_k}^2) \end{aligned} \tag{7}$$

where $k = 1, \dots, 30$ indexes the network, and $i, j = 1, \dots, 20$ indexes the nodes. X_k is a binary vector of length 30.

The prior distribution γ_0 , γ_1 , θ_0 , θ_1 , and ν had little to no impact on the parameter recovery of that parameter or on the parameter recovery of any other parameters in the model. The prior distributions of τ_Z and τ_Y did parameter recovery of τ_Z and τ_Y respectively; priors whose distributions that gave non-trivial weight to extremely large values resulted in more extreme values in the posterior distributions of τ_Z and τ_Y and worse parameter recovery.

Table 3. A summary of all priors considered in the sensitivity analysis. For each condition, the priors for a single parameter were varied and the other priors for fixed to the top row

γ_0	γ_1	τ_Z	θ_0	θ_1	ν	τ_Y
N(0,1)	N(0,1)	IG(5,0.5)	N(0,0.5)	N(0,1)	N(0,0.5)	IG(5,0.5)
N(0.5,0.5)	N(-0.6,0.6)	IG(10,1)	N(-0.2,0.2)	N(0.5,0.5)	N(-0.1,-0.1)	IG(10,1)
N(0,10)	N(0,10)	IG(1,1)	N(0,10)	N(0,10)	N(0,10)	IG(1,1)
N(0,100)	N(0,100)	IG(1,2)	N(0,100)	N(0,100)	N(0,100)	IG(1,2)

Similarly, prior distributions for τ_Z and τ_Y also had varying impacts on the posterior variance for other parameters. When IG(1,1) or IG(1,2) was used for τ_Z , the posterior variances for γ_0 and γ_1 were slightly larger, but when IG(1,1) or IG(1,2) was used for τ_Y , the posterior variances for θ_0 , θ_1 , and ν were between 3 and 5 times higher.

Despite some effects of the choice prior distribution, the posterior distribution of the mediation effect was generally robust to choice of prior, and only the posterior variance was impacted when extreme prior distributions were used for variance components. This does impact the utility of the model in that users must be thoughtful regarding their choice of prior distribution for the τ_Y .