Article

Journal of Educational and Behavioral Statistics
2019, Vol. 44, No. 2, pp. 210–240
DOI: 10.3102/1076998618814255
Article reuse guidelines: sagepub.com/journals-permissions

Modeling Social Networks as Mediators: A Mixed Membership Stochastic Blockmodel for Mediation

Tracy M. Sweet
Department of Human Development and Quantitative Methodology, University of Maryland

There are some educational interventions aimed at changing the ways in which individuals interact, and social networks are particularly useful for quantifying these changes. For many of these interventions, the ultimate goal is to change some outcome of interest such as teacher quality or student achievement, and social networks act as a natural mediator; the intervention changes the social networks of the teachers in schools, and teachers with certain types of social networks tend to use better teaching practices, for example. Due to lack of methodology, however, social networks have not been modeled as mediators. We present a new framework for modeling social networks as mediators in which a social network model is embedded into a mediation model and both models are estimated simultaneously. As a proof of concept, we introduce a new network model for mediation, applicable for interventions that affect subgroup structure. We provide a small simulation study to demonstrate the feasibility of this model and explore some potential operating characteristics. Finally, we apply our model to examine the effects of instructional coaches on teacher advice-seeking networks and subsequent changes in beliefs about mathematics.

Keywords: *evaluation; organization theory/change; research methodology; school/ teacher effectiveness; statistics*

1. Introduction

There are some educational interventions aimed at changing the ways in which individuals interact, whether it is an increase in collaboration among certain researchers, a push toward small learning communities of teachers, or breaking up adolescent cliques. Social networks may provide insight into not only how the intervention affected the ways individuals interact but also the mechanisms of how networks can affect individual outcomes. In some contexts, the network can even act as a mediator between the intervention and the outcome.

Of course, the ultimate purpose or theory of action may not be to change only the social or professional relationships among the individuals in the network, rather it is more likely that these interactions facilitate or influence individual outcomes. In educational studies, for example, changes in the social network are merely the mechanisms through which teachers become better teachers and students become better learners. For example, an initiative aimed at increasing teacher collaboration does so in the hopes of improving teacher quality and student outcomes.

Social network data in education have been used to measure access to professional resources such as expertise (Frank, Zhao, & Borman, 2004; Penuel et al., 2010), connect network structure with collective efficacy (Moolenaar, Sleegers, & Daly, 2012), link networks with formal organizational structure within schools (Spillane & Hopkins, 2013; Spillane, Kim, & Frank, 2012), and link network structure with top-down district policy (Coburn & Russell, 2008). Studying social networks of students has also been of interest; for example, networks have been used to investigate aggressive behavior and its relationship to peer networks (Espelage, Holt, & Henkel, 2003), to determine which factors are associated with new friendship ties (Frank, Muller, & Mueller, 2013), and to even link teacher attitudes about social behavior to student friendship networks (Gest & Rodkin, 2011). Social networks have even been considered as part of an intervention, for example, how networks affect intervention implementation (Daly, Moolenaar, Bolivar, & Burke, 2010; Moolenaar, Daly, & Sleegers, 2010) or how the intervention affects networks (Paluck, Shepherd, & Aronow, 2016).

Although social network data are being collected in education and psychological research, social network modeling has been largely absent in large-scale education studies. Further, we do not know of any studies in which networks have been used as mediators. We believe there are two reasons for this. First, until recently, there were very few statistical models that could easily accommodate multiple school network data inherent in these studies, and second, there are no known statistical social network models that accommodate a social network as a mediator.

The purpose of this article is therefore to introduce a general framework for modeling social networks as mediators. As an example, we will introduce the Hierarchical Mixed Membership Stochastic Blockmodel (HMMSB) for mediation, which builds on work by Airoldi, Blei, Fienberg, and Xing (2008); Sweet, Thomas, and Junker (2014); and Sweet and Zheng (2017). The article is organized as follows: We begin with an introduction to social network models focusing on network models for experimental interventions. Next, we provide an introduction to mediation analysis and describe how we can incorporate networks into a mediation model. We then introduce our framework, Hierarchical Network Models (HNMs) for mediation, for building models that accommodate networks as mediators and illustrate this with the HMMSB for mediation. We then provide a series of empirical analyses to illustrate how these models can be used in practice; we also include a real-world application using teacher advice-seeking network data. We conclude with some future directions for this work.

2. Social Network Models

The term *social network model* encompasses a wide variety of statistical models for modeling relational data. Some of these models consider the network or collection of network ties as the outcome, whereas others incorporate network data into the model to predict some other outcome. One issue with modeling social network data is that network ties are not mutually independent; that is, the likelihood that a tie exists from node i to node j is not independent of the other ties that may exist in the network. Intuitively, it is reasonable to assume that the probability of such a tie likely depends on whether there is a tie from j to i as well as the other ties i and j have with other nodes. Thus, models that assume independent observations should not be applied.

The purpose of this section is to first introduce the reader to social network models so that we can then augment them to model networks as mediators. Therefore, our overview will focus on social network models that will be used to build our network mediation models.

There are two broad classes for predicting social networks: models that specify the dependence structure among ties explicitly and models that accommodate tie dependence implicitly through the use of latent variables. The former class includes exponential random graph models (Robins, Pattison, Kalish, & Lusher, 2007; Wasserman & Pattison, 1996) in which the probability of observing a network is modeled as a function of network statistics and node covariates. The latter class is often called conditionally independent tie (or dyad) models in which each tie (or dyad) is modeled as independent conditional on some set of latent variables. Stochastic blockmodels (Anderson & Wasserman, 1992; Holland, Laskey, & Leinhardt, 1983), mixed membership stochastic blockmodels (MMSBs; Airoldi, Blei, Fienberg, & Xing, 2008), and latent space models (Hoff, Raftery, & Handcock, 2002) are examples of conditionally independent tie models.

We will focus on conditionally independent tie models for the purposes of this article, but these methods likely extend to exponential random graph models. We now define a conditionally independent tie model. Let A be an adjacency matrix where A_{ij} is the value of a tie from node i to node j. If these ties are binary, $A_{ij} = 1$ indicates the presence of a tie, and we write the model as

$$\begin{aligned} P(A|Z,X,\phi) &= \prod_{ij} P(A_{ij}|Z_{ij},X_{ij},\phi) \\ \operatorname{logit} P[A_{ij}=1|X_{ij},Z_{ij},\phi] &= f(\phi,X_{ij},Z_{ij}), \end{aligned} \tag{1}$$

where X is a collection of covariates at the node or dyad level, ϕ is the set of model parameters, and Z is a collection of latent variables. Thus, each tie A_{ij} is conditionally independent of every other tie given latent variables Z. The specific form the latent variables Z take depends on the model, although Z generally represents information about the individual nodes in the network.

Our proposed framework for modeling networks as mediators assumes that the network is the unit of observation; that is, we assume an intervention or experiment affects the entire network which in turn affects some outcome of interest. Thus, we require models that accommodate more than one network. Generalizing across some specific examples of multilevel models (e.g., Snijders & Kenny, 1999; Templin, Ho, Anderson, & Wasserman, 2003; Zijlstra, van Duijn, & Snijders, 2006), Sweet, Thomas, and Junker (2013) introduced a framework termed HNMs for modeling multiple networks. Specifically, this modeling framework assumes that networks are independent but not identically distributed as each network is generated by its specific network model.

Given an independent sample of networks $\mathbb{A} = (A_1, \ldots, A_K)$, covariates $\mathbb{X} = (X_1, \ldots, X_K)$, latent variables $\mathbb{Z} = (Z_1, \ldots, Z_K)$, and modeling parameters $\Phi = (\phi_1, \ldots, \phi_K)$, we can extend Equation 1 using the notation similar to that presented by Sweet et al. (2013). An HNM is given as

$$\begin{aligned} P(\mathbb{A}|\mathbb{X}, \mathbb{Z}, \Phi) &= \prod_k P(A_k|X_k, Z_k, \phi_k) \\ \prod_k P(A_k|X_k, Z_k, \phi_k) &= \prod_k \prod_{ij} P(A_{ijk}|Z_{ijk}, X_{ijk}, \phi_k) \\ \text{logit} P[A_{ijk} = 1|Z_{ijk}, X_{ij}, \phi_k] &= f(\phi_k, X_{ijk}, Z_{ijk}), \end{aligned} \tag{2}$$

where each ϕ_k is sampled from some super population. Similar hierarchical distributions can be specified for each network's set of latent variables Z_k as well.

Note that several multilevel extensions of conditionally independent tie network models have been introduced and versions of these models have been used in the education sciences. Sweet et al. (2013) introduced a multilevel extension to the latent space model which has been used to compare the effects of English-language learner (ELL) teachers versus non-ELL teachers on forming advice and instructional ties (Hopkins, Lowenhaupt, & Sweet, 2015), to examine whether beliefs about mathematics predict advice and instructional ties (Spillane, Hopkins, & Sweet, 2018), and to estimate the effect of proximity on workplace network ties (Spillane, Shirrell, & Sweet, 2017). Similarly, Sweet et al. (2014) and Sweet and Zheng (2018) extended the MMSB for multiple networks, and Sweet and Zheng (2018) presented an example relating teacher classroom management strategies with the level of integration among peer subgroups in that class.

3. Mediation Analysis

Mediation analysis refers to statistical models that accommodate a mediating variable, which is a variable involved in a causal process from independent variable to mediating variable and from mediating variable to dependent variable. An example of such a process could be an intervention on student social

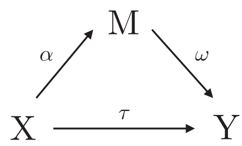

FIGURE 1. *A mediation model with independent variable* X*, dependent variable* Y*, and mediator* M*; causal paths from* X *to* M*, from* X *to* Y*, and from* M *to* Y *are denoted with an arrow. The regression coefficients* (α, ω, τ) *align with those given in Equation 3.*

skills (independent variable) which creates a sense of community (mediating variable) that subsequently improves class achievement (dependent variable).

The purpose of a mediation analysis is to determine the extent to which the relation between the independent variable and dependent variable is attributed to this third variable. The mediation analytic framework was first introduced by Baron and Kenny (1986) to properly disentangle the concept of mediation from that of moderation and has since grown to be not only an active area of research (e.g., Imai, Keele, Tingley, & Yamamoto, 2011; Krull & MacKinnon, 1999; MacKinnon & Fairchild, 2009; Yuan & MacKinnon, 2009) but also a standard topic for instruction (MacKinnon, 2008).

The single mediator model as introduced by Baron and Kenny (1986) is summarized in Figure 1. In this figure, X is the independent variable, M is the mediating variable, and Y is the dependent variable. This model can also be written as a series of three equations:

$$\begin{aligned} Y &\sim N(\beta_1 + \tau X, \sigma_1{}^2) \\ M &\sim N(\beta_2 + \alpha X, \sigma_2{}^2) \\ Y &\sim N(\beta_3 + \tau' X + \omega M, \sigma_3{}^2), \end{aligned} \tag{3}$$

where τ is often called the direct effect (of X on Y). The effect of the X on the mediating variable is α and the effect of M on Y is ω. The parameter τ' represents the effect of X on Y controlling for M.

Thus, if X predicts Y and M predicts Y, then when Y is regressed on both X and M, the effect of X is likely reduced when controlling for M. The mediated effect can be estimated by using either $\hat{\alpha}\hat{\omega}$ or $\hat{\tau} - \hat{\tau}'$; in an ordinary least squares (OLS) framework, these are mathematically equivalent (Yuan & MacKinnon, 2009).

According to the model given in Equation 3, any statistically significant association between X and M and M and Y can result in significant mediation effect. Moreover, associations among X, Y, and M are also found in models of moderation, and these are also present if M is a confounding variable. Thus, we need additional assumptions to prove that M is a mediating variable. The relationships among X and M and Y are causal; X causes M and M causes Y as shown in Figure 1. That these relation are causal defines mediation, and there are a variety of ways to assess whether these associations are causal. For the purposes of our article, we assume that these relationships are causal based on the social theory connecting our variables.[1]

Consider the effect of a professional development program for teachers on instructional quality. Teachers who participated in this program had higher ratings of instructional quality than teachers who did not participate. While it is certainly interesting that this program appears to improve instructional quality, one wonders why or how this program causes this improvement. A mediating variable in this example might be teacher efficacy. Suppose the program promoted building self-efficacy (stronger beliefs about their capabilities) so that participating teachers had higher levels of efficacy than nonparticipating teachers. Similarly, teachers with higher self-efficacy provide higher quality lessons. Thus, mediation models not only explain causal relations among a number of variables, they can also illuminate possible mechanisms to explain the reason behind the relationship between independent and dependent variables.

4. Modeling Networks as Mediators

Given the causal path from X to M to Y in Figure 1, there are a number of different ways to augment Equation 3 to accommodate a social network as a mediator. Before we introduce these, we present a running example to provide a context for these methods. Consider a school-based intervention such that participating in the intervention changes the way teachers in that school interact, and suppose the outcome of interest is also at the school level, such as a school-level measure of average teacher quality. In this context, the school is the unit of observation (or randomization in a randomized-controlled trial), so that X_k is an indicator for whether school k participates in the intervention and Y_k is the school-level outcome measure for school k. Then we propose using the school's network as the mediating variable M. It is important to note that since the unit of observation is the school (or network), our proposed methodology requires a collection of school networks; for the purposes of our example, assume there are K schools.

We consider two approaches to defining a network as a mediator. First, define the social network for school k using an adjacency matrix A_k, defined such that A_{ijk} is the value of a tie from individual i to individual j in school k.[2] The most obvious method of embedding A into M is to use a variable that can represent the social network as the mediating variable; that is, we substitute a network-level summary statistic of A_k for M (see top of Figure 2). A common network statistic is density, the proportion of observed ties out of all possible ties. In a binary network, we define density as $d = \frac{\sum_{i,j=1}^{n} A_{ij}}{n(n-1)}$, where n is the number of individuals in the network. Then our mediating variable is the network density for each school, $M = (d_1, d_2, \ldots, d_K)$. Returning to our example, we might assume that

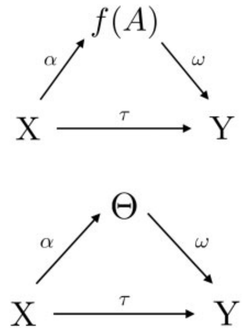

FIGURE 2. *Two methods of incorporating a network* A *into a mediation model are shown; use a summary statistic of the network by taking some function of the network* $f(\mathrm{A})$ *as the mediating variable (top) or fit a statistical social network model and use a parameter from that model as the mediator (bottom).*

our intervention increases the density of teacher networks which then improves teacher quality since teachers who are well-connected have access to more instructional resources.

One disadvantage with using a single network statistic is the fact that a single measure often fails to capture the full structure of the network. Similarly, a network summary statistic may be difficult to compare across networks, especially those that vary in size. On the other hand, if social theory is driving the analysis, using a single network statistic can be advantageous in that it can target a very specific aspect of the network. For example, if an intervention is hypothesized to increase collaboration among teachers across grades, the density of ties among teachers in different grades could be an important variable to include in the analysis. Another advantage of using a network statistic as a mediating variable is that standard statistical software can be used to fit these models.

The second approach, which is the focus and contribution of this article, is to use a parameter from a social network model for M (see bottom of Figure 2). Not only do social network models commonly account for the structure of the entire network (and its many interdependencies), these models can also identify network structures and social phenomena that may be difficult to capture with a network statistic. By embedding a social network model into a mediation model, we have not only created a new way to model networks as mediators but also an organizational framework with which to build mediation models for networks. We formally introduce this framework in the next section.

5. HNMs for Mediation

We propose a framework for modeling networks as mediators in which we fit a statistical social network model and use the parameter from that model as the mediator in a mediation model. Because we are interested in the effect of an intervention on a network, our unit of observation is at the network level, so social network models that accommodate multiple, independent networks are required. Therefore, our framework, HNMs for Mediation, extends the HNMs framework introduced by Sweet et al. (2013). To formally define an HNM for Mediation, let $A_1, \ldots, A_K$ represent a collection of independent networks such that A_k represents network k with a corresponding network model $P(A_k|\nu_k)$. Suppose our mediator is a parameter $\mathcal{M}$ such that $\mathcal{M}_k$ is contained in ν_k. We denote Y and X as the associated dependent and independent variables, and α, ω, and τ' are the causal effects defined in Equation 3. Then, a general form for an HNM for Mediation is given as

$$\begin{aligned} P(A|X,\nu) &= \prod_{k=1}^{K} P(A_k|\nu_k) \\ \mathcal{M}_k &\sim F(\mathcal{M}_k|X_k,\alpha) \\ Y_k &\sim G(Y_k|\mathcal{M}_k, X_k, \omega, \tau'), \end{aligned} \tag{4}$$

where G and F are generic distributional assumptions. To better align with Equation 3, we can specify $\mathcal{M}_k$ in Equation 4 as having a normal distribution; that is, we could specify $\mathcal{M}_k \sim N(\beta_2 + \alpha X_k, \sigma_M^2)$. Similarly, we can write $Y_k \sim N(\beta_3 + \tau' X_k + \omega \mathcal{M}_k, \sigma_Y^2)$.

Note that Equation 4 is purposively general to highlight the variability in models that can be built from this framework, in terms of the network model, the mediating network model parameter as well as how the mediation model is specified. That said, we are not claiming that any model fitting this form is a feasible mediation model, only that this is a general way of specifying a mediation model and that network mediation models can be written in this form.

To better understand Equation 4, let us return to our school-based intervention such that X is the treatment condition indicator. Then we believe that something about the network changes due to participating in the experiment, and this something can be measured by $\mathcal{M}$. Finally, we also hypothesize that $\mathcal{M}$ differentially impacts some school-level variable such as average teacher quality. We further specify this example in the next section.

5.1. HMMSBs for Mediation

As a proof of concept, we introduce the first HNM for Mediation and call this model the HMMSB for Mediation; the long name points to the various components used in building the model. We will briefly introduce these components before we present the full model.

Stochastic blockmodels are social network models that represent the probability of observing a network tie between two individuals as a function of their respective block memberships; the common assumption is that if the two individuals are in the same block or group, they are more likely to share a tie than if they are in different blocks. These models are particularly useful for networks with some kind of subgroup structure such as friendship ties among adolescents or workplace ties among teachers. A simple stochastic blockmodel (for a binary network) is given as

$$\begin{aligned} A_{ij} &\sim \text{Bernoulli } g_i^T B g_j \\ g_i &\sim \text{Multinomial } (p) \\ B_{\ell m} &\sim \text{Beta } (a_{\ell m}, b_{\ell m}), \end{aligned} \tag{5}$$

where A_{ij} is the indicator of a tie from i to j, and g_i and g_j are sampled from a multinomial distribution and denote the block memberships for i and j, respectively. The parameter B represents a block–block tie probability matrix such that $B_{\ell m}$ is the probability of a tie from a node in block ℓ to a node in block m. Prior distributions for the elements in B can be specified such that within-block ties are more likely than between-block ties; that is, the shape parameters $a_{\ell m}$ and $b_{\ell m}$ are chosen to place more probability on values near 0 or near 1.

Mixed membership refers more generally to larger class of mixture models in which group assignment is not fixed but instead parameterized by a probability vector (Airoldi, Blei, Erosheva, & Fienberg, 2014). The concept of mixed membership has been applied in a number of statistical models across a wide variety of disciplines ranging from science publications (Erosheva, Fienberg, & Lafferty, 2004) to predicting trends in disability (Manrique-Vallier, 2014) to problem-solving strategies (Galyardt, 2012).

MMSBs were originally introduced by Airoldi et al. (2008) as an extension of stochastic blockmodels. Block membership is no longer limited to a single block; each node belongs to all blocks with nonzero probability. An MMSB can be written as

$$\begin{aligned} A_{ij} &\sim \text{Bernoulli } ({S_{ij}}^T B R_{ji}) \\ S_{ij} &\sim \text{Multinomial } (\theta_i) \\ R_{ji} &\sim \text{Multinomial } (\theta_j) \\ \theta_i &\sim \text{Dirichlet } (\xi\gamma) \\ B_{\ell m} &\sim \text{Beta } (a_{\ell m}, b_{\ell m}), \end{aligned} \tag{6}$$

where θ_i represents the probability that node i belongs to each block. This membership vector θ also determines the block membership for each interaction from i to j; that is, θ_i affects the distributions of S_{ij} and R_{ji}, the vectors indicating the block membership of i and j, respectively, when i is sending a tie to j. Thus, the probability of a tie from i to j depends on both the specific membership for that interaction as well as B. Note that the distribution for θ_i in Equation 6 differs from that in Airoldi et al. (2008); we use the paramerization proposed by Erosheva (2003) such that ξ determines the relative center of the Dirichlet distribution and γ determines the shape of the Dirichlet distribution. In practice, we can think of ξ as a vector that shows the relative probability of belonging to each subgroup such that equal entries of ξ generally result in similar subgroup sizes. On the other hand, γ determines how extreme the membership probability is; small values of γ result in nodes belonging to only one block with probability near 1 and the other blocks with probability near 0.

We now consider the following extensions to the MMSB. First, Sweet et al. (2014) introduced a multilevel extension of the MMSB termed the HMMSB. One way to specify this model is given as

$$\begin{aligned} A_{ijk} &\sim \text{Bernoulli}\,(S_{ijk}{}^{T} B R_{jik}) \\ S_{ijk} &\sim \text{Multinomial}\,(\theta_{ik}) \\ R_{jik} &\sim \text{Multinomial}\,(\theta_{jk}) \\ \theta_{ik} &\sim \text{Dirichlet}\,(\xi\gamma_{k}) \\ \gamma_{k} &\sim \text{Gamma}\,(c, d) \\ B_{\ell mk} &\sim \text{Beta}\,(a_{\ell mk}, b_{\ell mk}), \end{aligned} \tag{7}$$

where k indexes the network. Not only do these multilevel extensions now accommodate multiple, independent networks, but they also estimate a network-level (network-specific) parameter γ_k.

Further, Sweet and Zheng (2017) proposed γ as a measure of subgroup integration. Networks generated from MMSBs with values of $\gamma < 0.1$ tend to have very insular subgroups, whereas networks generated from models with $\gamma > 0.5$ have very integrated networks.[3] Sweet and Zheng (2018) then proposed a new model in which γ for each network is regressed on some network-level covariate. This model is the same as Equation 7 except for the distribution of γ which is given as

$$\begin{aligned} \gamma_k &\sim \exp(\beta' X_k) \\ \beta &\sim N(c, d), \end{aligned} \tag{8}$$

where X is a network-level covariate.

Note that in Equation 8, we assume that the value of γ for each network—that is, the measure of subgroup integration for that network—covaries with some network attribute. For example, consider elementary school teacher advice-seeking networks. These tend to have subgroup structure since teachers tend to cluster based on grade assignment. However, there is variability in how insular these clusters might be. We might relate subgroup insularity of these teachers with the vision of the administrative leaders. If the administration believes that vertical alignment is a priority (X in Equation 8), they likely will take steps to have teachers interact with teachers who do not teach the same grade, resulting in less insular teacher networks.

Moreover, we argue Equation 8 could be also used in the context of an intervention such that X is a binary or categorical variable indicating assignment to an experimental condition. It is also possible that X is a continuous variable measuring some kind of intervention dosage. We then look at the effects of such an intervention on subgroup insularity γ.

We now present the HMMSB for Mediation. Let X be an intervention that affects the level of subgroup insularity γ which in turn affects some other network-level outcome Y. The full mediation model is given as

$$
\begin{aligned}
A_{ijk} &\sim \text{Bernoulli}\left(S_{ijk}{}^{T} B R_{jik}\right) \\
S_{ijk} &\sim \text{Multinomial}\left(\theta_{ik}\right) \\
R_{jik} &\sim \text{Multinomial}\left(\theta_{jk}\right) \\
\theta_{ik} &\sim \text{Dirichlet}\left(\xi\gamma_k\right) \\
B_{\ell mk} &\sim \text{Beta}\left(a_{\ell mk}, b_{\ell mk}\right) \\
\gamma_k &\sim N\left(\beta_{01} + \alpha X_k, \sigma_1{}^2\right), \gamma_k > 0 \\
Y_k &\sim N\left(\beta_{02} + \tau' X_k + \omega\gamma_k, \sigma_2{}^2\right) \\
\beta_{01} &\sim N\left(\mu_{\beta_{01}}, \sigma_{01}{}^2\right) \\
\beta_{02} &\sim N\left(\mu_{02}, \sigma_{02}{}^2\right) \\
\alpha &\sim N\left(\mu_\alpha, \sigma_\alpha{}^2\right) \\
\tau' &\sim N\left(\mu_{\tau'}, \sigma_{\tau'}{}^2\right) \\
\omega &\sim N\left(\mu_\omega, \sigma_\omega{}^2\right) \\
\sigma_1{}^2 &\sim \text{Inv} - \text{Gamma}(a, b) \\
\sigma_2{}^2 &\sim \text{Inv} - \text{Gamma}(c, d),
\end{aligned}
\tag{9}
$$

where α, τ', and ω are defined as in Equation 3; α is the effect of the intervention on the network; τ' is the conditional effect of intervention on the outcome variable; and ω is the conditional effect of the network on the outcome. As in Equation 3, we define the mediated effect as $\alpha\omega$. First, note that our distribution for γ differs from the distribution proposed by Sweet and Zheng (2018) in that we use a truncated normal distribution. We found that a truncated distribution generated more appropriate values of γ given a nominal covariate X; Sweet and Zheng (2018) used a continuous predictor in their model.

Note also that in Equation 9, we include prior distributions for the entries of B_k across all networks, but in practice, we do not recommend estimating B. In fact, Sweet and Zheng (2017) discussed an identifiability issue between the values of B and γ and suggested that B remained fixed across networks to optimize the precision in estimating each γ_k. This also means that γ is a relative parameter; each γ_k depends on that network's values of B so the value for each γ_k is not meaningful when taken alone.

Returning to our running example, suppose now that our intervention X is a professional development program that promotes teacher collaboration across grades, subjects, and levels of experience. As a result, teachers who participate in this program are more likely to form professional relationships with teachers outside of their instructional area than teachers who do not. Thus, treated networks have larger values of γ. Further, integrated networks result in larger spans of resource sharing and improved instructional quality Y. Therefore, instructional quality is predicted by both participating in the professional development X and having integrated subgroups γ.

5.2. Estimation

To estimate the parameters in Equation 9, we developed an MCMC algorithm (Gelman, Carlin, Stern, & Rubin, 2013) in which parameters are updated using Gibbs, Metropolis, or Metropolis–Hastings (M-H) updates and estimated using R (R Core Team, 2018). The joint likelihood is given as

$$\begin{aligned}
&P(\omega,\alpha,\tau',\sigma_1^2,\sigma_2^2,\theta,S,R,\xi,B,X,Y,A) \\
&= \prod_k \prod_{i \neq j} \Big(P(A_{ijk}|S_{ijk},B_k,R_{jik}) P(S_{ijk}|\theta_{ik}) P(R_{jik}|\theta_{jk}) \Big) \prod_k \prod_i \Big(P(\theta_{ik}|\xi_k,\gamma_k) \Big) \\
&\quad \prod_k P(Y_k|X_k,\beta_{02},\tau',\gamma_k,\omega,\sigma^2{}_2) \prod_k \Big(P(\gamma_k|X_k,\alpha,\beta_{01},\sigma^2{}_1) P(B_k) P(\xi_k) \Big) \\
&\quad P(\beta_{01}) P(\beta_{02}) P(\alpha) P(\omega) P(\tau') P(\sigma^2{}_1) P(\sigma^2{}_2).
\end{aligned} \tag{10}$$

The parameters θ, S, R, and B can be updated using Gibbs updates. The complete conditionals for each parameter in network k are

$$\begin{aligned}
P(\theta_{ik}|\ldots) &\propto \text{ Dirichlet } \left(\xi_k \gamma_k + \sum_j S_{ijk} + \sum_j R_{ijk} \right) \\
P(S_{ijk}|\ldots) &\propto \text{ Multinomial } (p) \\
& \quad p_h = \theta_{ikh} B_{h\ell^\star}{}^{A_{ijk}} (1 - B_{h\ell^\star})^{(1-A_{ijk})} \\
P(R_{ji}|\ldots) &\propto \text{ Multinomial } (q) \\
& \quad q_h = \theta_{ikh} B_{m^\star h}{}^{A_{ijk}} (1 - B_{m^\star h})^{(1-A_{ijk})} \\
B_{\ell mk} &\propto \text{ Beta } \Big(a_{\ell k} + \sum_{(ijk)^\star} A_{ijk}, b_{mk} + \sum_{(ijk)^\star} A_{ijk} \Big),
\end{aligned} \tag{11}$$

where $l^\star$ and $m^\star$ are the group membership indicated by R_{jik} and S_{ijk}, respectively, and $(ijk)^\star$ is a specific subset of (i,j) in network k such that S_{ijk} indicates block ℓ and R_{jik} indicates block m.

The remaining parameters require Metropolis or M-H updates. To update ξ_k for a given network, we use a Dirichlet proposal distribution centered at the current value of ξ_k. Then, we propose $\xi_k{}^{(s+1)} \sim$ Dirichlet $(u_{\xi_k} g \xi_k{}^{(s)})$, where u_{ξ_k}

is a tuning parameter and g is the number of blocks, which is chosen a priori. Similarly, γ_k can be updated using the following proposal distribution $\Gamma(u_{\gamma_k}, \frac{u_{\gamma_k}}{\gamma_k})$ such that u_{γ_k} is a tuning parameter.

For parameters with normal priors, β_{01}, β_{02}, α, τ', and ω, we update using random walk M-H updates. Finally, $\sigma^2{}_1$ is updated using a Metropolis update using an Inv-Gamma $\left(\frac{u_\sigma^2}{\sigma_1^2}, u_\sigma^2\right)$ proposal distribution with tuning parameter u_σ^2, and σ_2^2 is updated via Gibbs with an Inv-Gamma complete conditional with hyperparameters given as $\left(c + \frac{K}{2}, d + \frac{1}{2}\sum_k (Y_k - (\beta_{02} + \tau' X_k + \omega\gamma_k))^2\right)$.

The MCMC algorithm summarizes how one would fit a HMMSB for Mediation in theory. In practice, we require several constraints. Sweet and Zheng (2017) recommend fixing B along with specifying the number of blocks to be the same across networks. They found that γ is actually a relative measure and that identifiability issues exist when B and the number of blocks vary across networks. They also found that γ is best recovered when γ is between 0.05 and 1 so we directly constrain $\gamma > 0$ through the use of a truncated normal distribution (see Equation 9).

Finally, we discuss our choice for estimating the HMMSB for Mediation as one large Bayesian model as opposed to estimating the HMMSB in two stages. First, our measure of subgroup insularity is a parameter from the MMSB, so there is estimation error associated with γ_k for each network. We need to include that error in our mediation model so that the posterior variance for our mediation effect is not misleadingly small, which would increase type I error rates as well as overinflate power. Second, estimating the full mediation model is not much more computationally expensive than estimating the MMSB for each model and then using OLS to estimate α, τ', and ω. Given the advantages of incorporating estimation error for γ into the model, it is worth the minimal computational cost to include these parameters in the MCMC.[4]

6. Empirical Analyses

6.1. Simulation Studies

We present a small simulation study to demonstrate fitting the HMMSB for Mediation in practice. We simulate network and outcome data for a variety of network sizes, numbers of networks, and various mediation and direct effects. A summary of the simulated data sets is given in Table 1.

For each simulated data set, we construct a binary variable X such that exactly half of the networks are treated ($X = 1$). We fix the number of blocks in each network to be three, and for simplicity, we fix $\xi = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, which generates three subgroups that are approximately equal in size. We then use the following data generating model:

TABLE 1.
Summary of Simulated Data Sets

Simulated Data Set	Number of Networks	Network Size	τ'	ω
1	20	15	.5	1
2	20	30	.5	1
3	40	15	.5	1
4	40	30	.5	1
5	20	30	0	0.5
6	20	30	0	1
7	20	30	.5	0.5

Note. All data sets have $\alpha = .5$.

$$\begin{aligned}
\gamma_k &\sim N(0.05 + \alpha X, 0.25), \quad \gamma_k > 0, \quad k = 1, \ldots, K \\
Y_k &\sim N(\tau' X_k + \omega\gamma_k, 0.25) \\
\theta_{ik} &\sim \text{Dirichlet}\ (\gamma_k \xi), \quad i = 1, \ldots, n \\
S_{ijk} &\sim \text{Multinomial}\ (\theta_{ik}), \quad i,j = 1, \ldots, n \\
R_{jik} &\sim \text{Multinomial}\ (\theta_{jk}), \quad i,j = 1, \ldots, n \\
A_{ijk} &\sim \text{Bernoulli}\ ({S_{ijk}}^T B R_{jik}),
\end{aligned} \tag{12}$$

such that K is the number of networks and i,j index is the n nodes in each network. In addition, we fix B to be a 3×3 matrix with diagonal entries of 0.6 and off-diagonal entries of 0.005.

To better understand the data generating model, we plot the simulated networks from Simulation 1 (Table 1) for which there were 20 networks of 15 nodes each. Figure 3 shows the effect of $\alpha = .5$ on network subgroup structure; the 10 control networks (top two rows of networks) have more insular subgroups than the 10 treated networks (bottom two rows of networks). Specifically, γ_k for the treated networks are generally larger than the γ_k for the control networks. Recall that small values of γ generate block membership probabilities such that most of the probability mass is in a single block and larger values of γ generate block probabilities that are less extreme. As a result, nodes in the treated networks are more likely to belong to multiple blocks, and the networks that result appear to have more integrated networks. In Figure 3's treated networks, some of the subgroups are so integrated that the three blocks are no longer distinguishable.

Simulation 1 also assumes a positive effect of the mediator on the outcome variable. Figure 4 shows a scatterplot of the outcome Y versus γ; the positive association is visually apparent when $\omega = 1$.

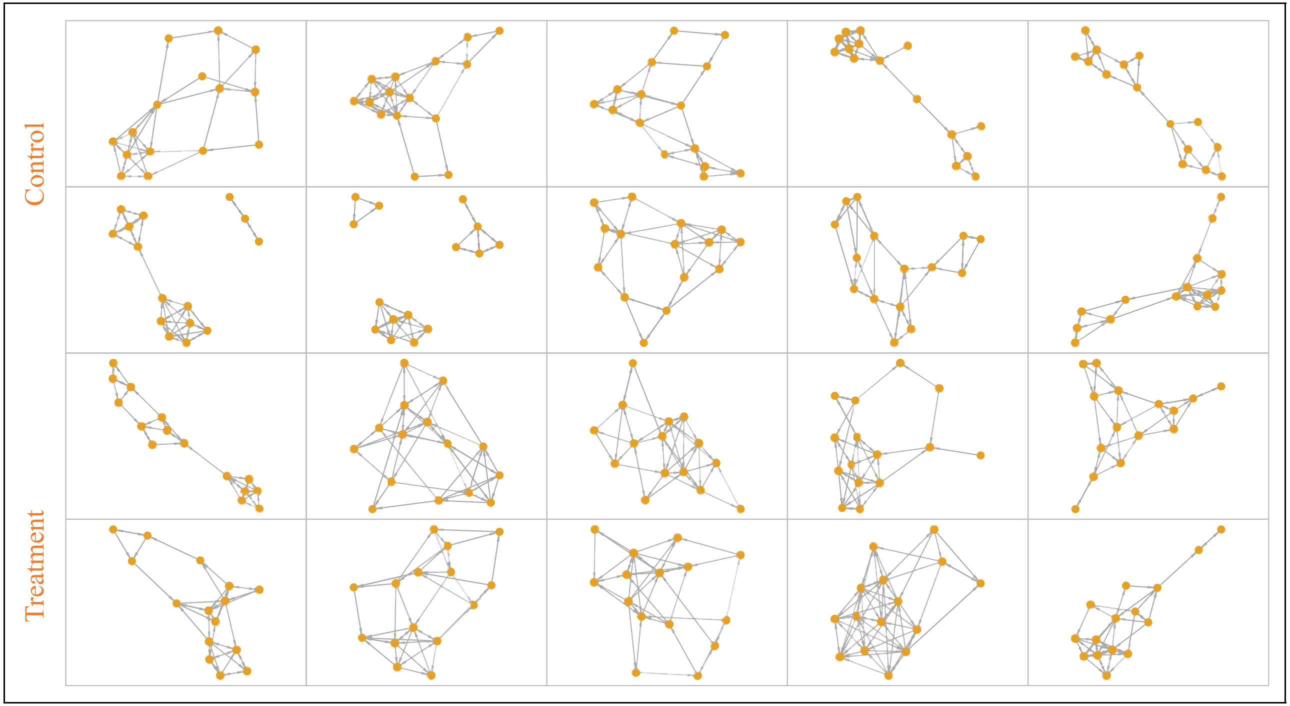

FIGURE 3. *Networks simulated from a Hierarchical Mixed Membership Stochastic Blockmodel for Mediation in which the intervention increases subgroup integration.*

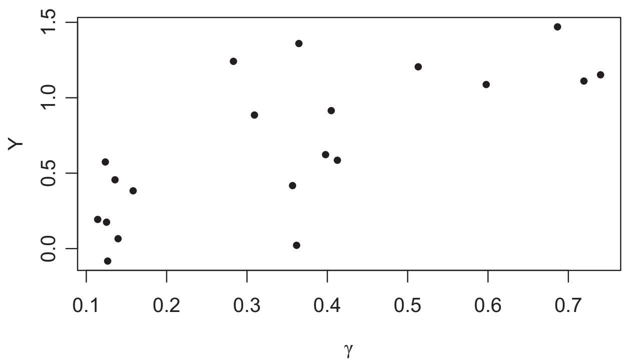

FIGURE 4. *Plot of the simulated outcome* Y *versus* γ *shows a positive effect of the mediator on the outcome variable.*

For each simulation setting described in Table 1, we generated 50 data sets[5] and fit the HMMSB for Mediation given in Equation 9. The priors used for each model fit are given as

$$\begin{aligned}\beta_{01} &\sim N(0, 0.01)\\ \beta_{02} &\sim N(0, 9)\\ \alpha &\sim N(0, 9)\\ \tau' &\sim N(0, 9)\\ \omega &\sim N(0, 9)\\ \sigma_1^2 &\sim \text{Inv} - \text{Gamma}(10, 0.5)\\ \sigma_2^2 &\sim \text{Inv} - \text{Gamma}(6, 2).\end{aligned} \tag{13}$$

There are some important things to note regarding how one specifies prior distributions. Because prior distributions are often seen as subjective, uninformative priors are preferred in the absence of prior knowledge. However, using prior distributions that are flat does not necessary achieve the goal of a non-informative prior. Consider, for example, the distribution of γ, which is best estimated when γ is between 0 and 1 (Sweet & Zheng, 2017). Then, a prior distribution on β_{01} or α that is completely flat is actually very informative as β_{01} should be quite small and α should also be less than 1. Our aim was to choose weakly informative prior distributions; the probability distribution for each prior covers what we considered to be a reasonable range of parameter values. We further explore the effects of prior distribution specification in Section 6.1.1.

Because the purpose of this simulation study is to show that the HMMSB for Mediation is both feasible and useful, we examine both parameter recovery of the

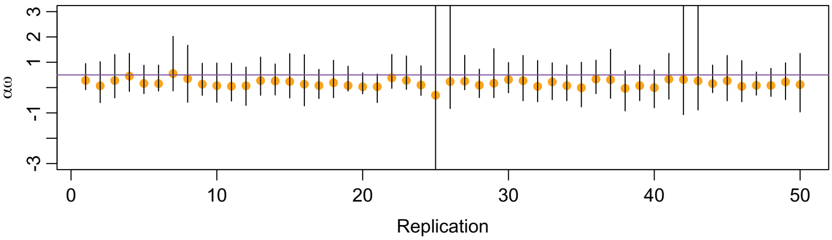

FIGURE 5. *Posterior summaries of the mediated effect αω from 50 simulated data sets each with 20 networks and 15 nodes; maximum a posteriori (MAP) estimates are plotted as points and segments represent the 95% equal-tailed credible intervals.*

mediation effect (αω) and how the network size, number of networks, and oracle values of the parameters affect estimation and recovery.

We begin by discussing the first simulation set consisting of 20 networks, each with 15 nodes, and a mediation effect of $(0.5)(1) = 0.5$. For each of the 50 replications, we ran our MCMC sampler twice using overdispersed starting points, and we examined traceplots to assess convergence and calculated autocorrelation to determine optimal thinning. Tossing the first 2,000 samples and retaining every 150th subsequent draw, we have posterior samples of size 374. The posterior mode and 95% equal-tailed credible interval for the mediated effect (αω) is shown in Figure 5.

There are several things to note about Figure 5. First, the credible intervals cover the true value of the mediated effect for all 50 replications. Second, there are some instances where the credible intervals are extremely wide. In these replications (e.g., replications 25, 26, 42, and 43), we found that the posterior distribution of α had extremely high variance. In retrospect, we should have not chosen such an informatively weak prior distribution for α since this prior allows values of α much larger than 1.

Figure 6 shows posterior summary information for the mediated effect (αω) for all seven simulations. The scale of the y-axis is the same across all plots so that we can easily compare them. As the network size increases (while the number of subgroups is fixed, effectively increasing block size), we find more precise estimates of the mediation effect and it appears that the posterior modes become more accurate. This makes sense since large block sizes generally improve our ability to estimate γ. Posterior variances similarly decrease when we increase the number of networks; the number of networks is effectively our sample size. Thus, posterior variance is minimized when both the number of networks and the network size (or block size) increase. As a result, we find that in Simulations 2 through 7, we have fewer replications with extremely large posterior variances than we had in Simulation 1.

Regarding the effect of τ', we fixed the number of networks and network size to 20 and 30, respectively, and compared $\tau' = 0$ with $\tau' = 0.5$. Figure 6 shows

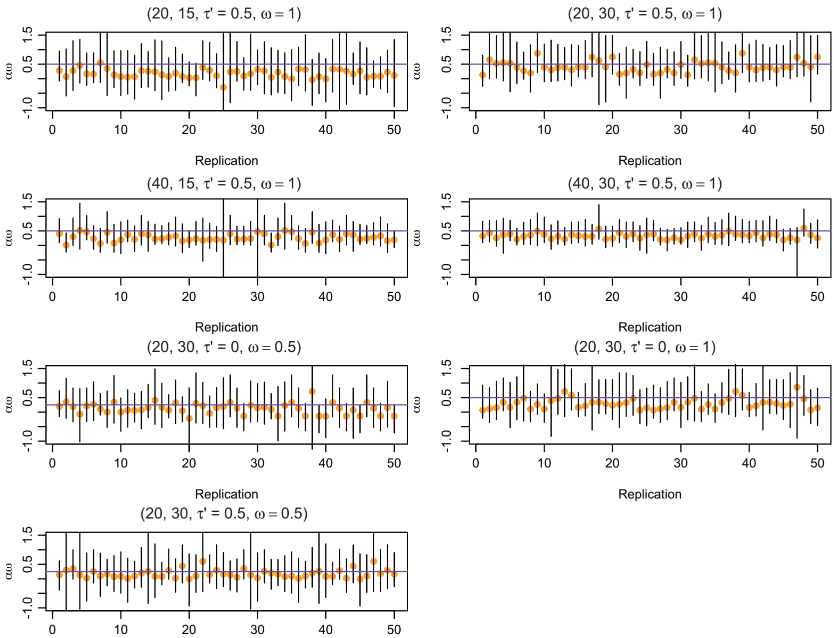

FIGURE 6. *MAP estimates and 95% equal-tailed credible intervals for all seven simulations for various numbers of networks, numbers of nodes, and values of τ and ω; $\alpha = .5$ for all simulations (see Equation 9).*

very little effect of the value of τ' on estimating $\alpha\omega$; that is, full mediation and partial mediation do not differ in terms of mediating parameter recovery.

6.1.1. Prior sensitivity analysis. The simulation study suggests that how one specifies prior distributions may affect posterior samples of the mediated effect. Recall in Figure 6 that larger sample sizes resulted in fewer replications with large posterior variance, which suggests that our choice of prior distribution for Simulation 1 may have been too informative. In our attempt to use a weak prior distribution for α of $N(0, 9)$, we essentially implied that we believed α would likely be between -3 and 3 and that values as small as -6 or as large as 6 would also be plausible, which was actually very informative.

Therefore, we present a small sensitivity study to explore how different prior distributions impact the posterior distribution of the mediated effect $\alpha\omega$. Table 2 presents the prior distribution for Simulation 1 along with four other conditions. Because α and β_{01} largely determine γ as well as σ_1^2, we first explore how each prior distribution specification affects the posterior distribution for $\alpha\omega$. We tighten the prior distribution so that $|\alpha| < 2$ with probability near 1, followed by a less informative prior for σ_1^2 and a less informative prior on β_{01}. Finally, we

TABLE 2.
Specified Prior Distributions for Sensitivity Analysis

Prior Specification	β_{01}	α	σ_1^2	$\beta_{02}, \tau', \omega$	σ_2^2
Simulation 1	$N(0, 0.01)$	$N(0, 0.5)$	$\Gamma^{-1}(10, 0.5)$	$N(0, 9)$	$\Gamma^{-1}(6, 2)$
Stronger prior on α	$N(0, 0.01)$	$N(0, 0.25)$	$\Gamma^{-1}(10, 0.5)$	$N(0, 9)$	$\Gamma^{-1}(6, 2)$
Weaker prior on σ_1^2	$N(0, 0.01)$	$N(0, 0.25)$	$\Gamma^{-1}(5, 1)$	$N(0, 9)$	$\Gamma^{-1}(6, 2)$
Weaker prior on β_{01}	$N(0, 0.5)$	$N(0, 0.25)$	$\Gamma^{-1}(10, 0.5)$	$N(0, 9)$	$\Gamma^{-1}(6, 2)$
Weaker priors on β_{02}, τ', ω, and σ_2^2	$N(0, 0.01)$	$N(0, 0.25)$	$\Gamma^{-1}(10, 0.5)$	$N(0, 100)$	$\Gamma^{-1}(3, 1)$

explore weaker priors on the parameters in outcome model: β_{02}, τ', ω, and σ_2^2. For each row in Table 2, we use the same 50 data sets used in Simulation 1 and fit models using the prior distributions indicated.

Figure 7 shows the 95% equal-tailed credible intervals and posterior modes for each set of prior distributions. With a more appropriate prior on α (Figure 7, top right), we no longer see extremely large posterior variances on $\alpha\omega$ although some simulations have larger variance than others. Note that the prior distribution $\alpha \sim N(0, 0.25)$ is not extremely informative; the distribution is not centered at the truth but simply includes our prior belief that α is between -1 and 1. A weaker prior on σ_1^2 has little effect on the posterior variance of $\alpha\omega$, and as the prior distribution on β_{01} accommodates the belief for larger absolute values of β_{01}, we find very slight decreases in the posterior variances of $\alpha\omega$, due to an increased posterior variances of β_{01}. Finally, even rather large changes in the priors of the other parameters, β_{02}, τ', ω, and σ_2^2, have very little effect on the posterior distribution of $\alpha\omega$.

These results are mostly comforting in that model estimation for the HMMSB for Mediation is impacted very little by priors. The only exception is using informatively weak or incorrect priors for α, and this appears to be an issue only when network size is small ($n = 15$).

6.1.2. Misspecified models. We also briefly explore whether mediation effects can be recovered when different parts of the HMMSB for Mediation are misspecified. We consider two examples, one in which we expect a small impact and one in which we expect a large impact.

The first misspecification example involves the number of blocks or subgroups, which is selected by the user. In the examples above in Section 6.1, our data generating model and our fitted model both used three blocks. In this example, our data generating model is similar but assumes there are six blocks in each network. The fitted model still assumes three blocks.

We chose this example to highlight the ability of the HMMSB for Mediation to estimate the mediated effect when the number of blocks is misspecified. As

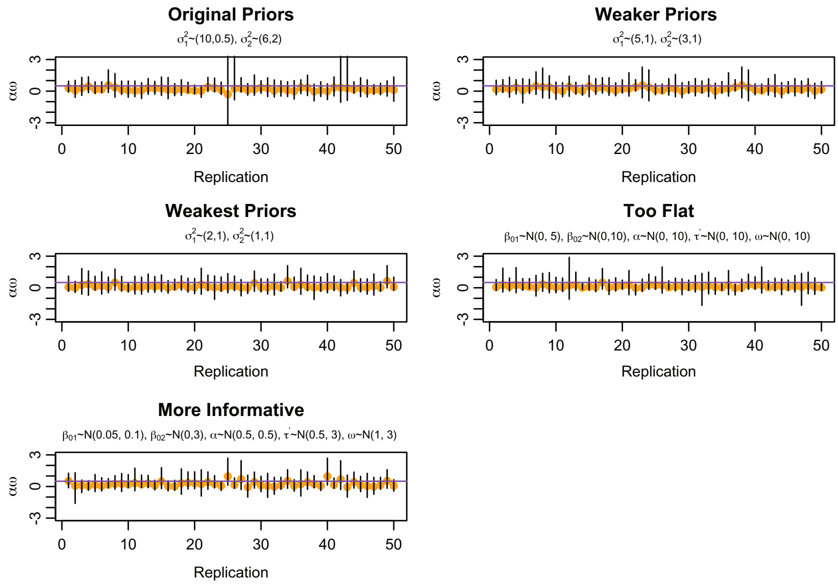

FIGURE 7. *MAP estimates and 95% equal-tailed credible intervals for the mediated effect αω under a variety of prior distribution specifications suggest that the prior specification for α can impact the posterior distribution for αω but priors for other parameters have little or no impact.*

discussed by Sweet and Zheng (2017), there is an identifiability issue among the number of blocks, the entries of the block–block tie probability matrix B and γ. Since we fix B when estimating the model and the fix the number of blocks to be the same across networks (even though it is the incorrect number of blocks), γ_k in one network is still estimated accurately relative to the other networks.

The posterior samples for the mediated effects $\alpha\omega$ are given in Figure 8 (left). The posterior samples recover the true effect with similar posterior variance and posterior modes.

The second example is a much stronger misspecification. In this model, the value of γ no longer affects the subgroup structure of the network and in fact is not a mediator in the model at all. The data generating model for the network is a multilevel latent space model (see Hoff et al., 2002; Sweet, Thomas, & Junker, 2013) such that treated networks are more dense than control networks. The outcome variable is therefore a function of the intervention and a randomly generated γ, so that γ is independent of both X and Y.

Thus, in this example, γ no longer mediates the relationship between the intervention and the outcome and this is evident in the figure summarizing the posterior distribution for $\alpha\omega$ in Figure 8 (right). As expected, we find no evidence of an effect of the intervention on the network subgroup insularity, and we find

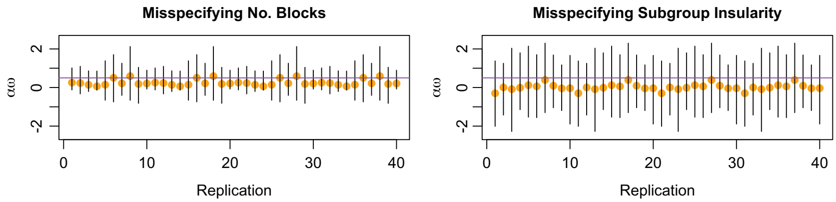

FIGURE 8. *Simulations in which the model was misspecified. The number of blocks in the generative model was more than in the fitted model (left) and the generative model did not account for* γ *as the network was generated from a different network model (right).*

no evidence of an association between subgroup insularity and the outcome. The posterior modes are all very close to 0 and the posterior variance in each replication suggests that each γ_k was estimated with more error than usual; this is unsurprising as we purposely generated networks without subgroup structure.

These examples further illuminate how the HMMSB for Mediation would behave in practice. Under user misspecification, such as the number of blocks, the model is still able to recover the true mediated effects which is important since one rarely knows the number of blocks to choose a priori. Further, when subgroup insularity is not a mediator, the estimated mediated effect is near 0, so the model is not finding effects that do not exist.

6.2. Coaches and Mathematics Beliefs

We present an application examining whether teacher advice-seeking networks mediate the effect of introducing coaches into schools on changes in beliefs about instruction. We use data from elementary schools in a school district that come from a series of studies on distributed leadership and have been used in education (Spillane & Hopkins, 2013; Spillane, Hopkins, & Sweet, 2015, 2018).[6] These data contain survey items collected from staff members in 14 schools regarding teaching practices and beliefs as well as advice-seeking nominations collected each spring over the course of 5 years (2010–2013, 2015).

Throughout the data collection period, eight schools received mathematics coaches to support and improve mathematics instruction. Although we do not have data on the quality of instruction, we do have data on beliefs about mathematics instruction. We expect that beliefs about instruction move in a more positive direction in schools with math coaches. One reason may be that coaches have resources and when presented with these new resources, teachers change the way they think about mathematics instruction. However, there is also evidence that instructional coaches are not only providing instructional support but are also poised to act as brokers within and across schools (Hopkins, Ozimek, & Sweet, 2017). Thus, we are interested in how the presence of an instructional coach changes network structure and whether that leads to changes in teacher outcomes.

While we do not have a randomized intervention, we can still use a HMMSB for Mediation to model the effect of introducing a math coach to a school. There are eight schools who received mathematics coaches: two schools were assigned a coach in 2011, five additional schools received coaches in 2012, and one final school received their coach in 2013. We are interested in the effects that these coaches had on the structure of the mathematics advice-seeking networks and whether these network structures mediated changes in beliefs about mathematics. The eight schools in the year they were assigned a coach are considered our treated schools; the eight control schools are schools of similar size from the same year who did not receive a coach.

We have two measures of beliefs about mathematics; both are constructed from a set of survey items. One measure targets beliefs about how focused mathematics instruction should be on procedures and skills (vs. inquiry). The other measure concerns beliefs about how student centered (vs. teacher centered) the classroom should be. Because we are interested in whether a coach caused these beliefs to change, our outcome variables will be denoted *change in procedural beliefs* and *change in student-centered beliefs*. For each school, we compare the belief (procedural or student centered) collected in the end of the year prior to receiving a coach with the belief collected after having had a coach. We do the same for our control networks. For example, for the two schools that received coaches in 2011, we compared beliefs in 2010 with beliefs in 2011 and then compared beliefs in 2010 and 2011 for two control schools.

For social networks, we use survey items in which teachers nominated those to whom they sought information and advice regarding mathematics instruction throughout the school year. These nominations were used to construct binary, directed networks for each school; we also assume that schools constitute separate networks and ignore the very small percentage of ties that occurred across schools.

Figure 9 shows the mathematics advice-seeking networks in the eight schools in each condition. The networks in the top two rows come from schools in their first year of having a mathematics coach and the bottom two rows show the matched networks who did not receive a coach. What may at first seem surprising is the number of isolated nodes in each network. These are nodes that responded to the survey but did not nominate any teachers to whom they sought advice regarding mathematics instruction, nor were they nominated as providing such advice. This is not surprising, given that these staff members work in elementary schools; the advice-seeking networks around language arts had fewer isolated nodes and generally more ties. Further, there appear to be more isolated nodes in the treated networks than the control networks. This may in part be due to how those schools were selected for coach assignment; these schools may have had lower math achievement scores or lower fidelity in general regarding district

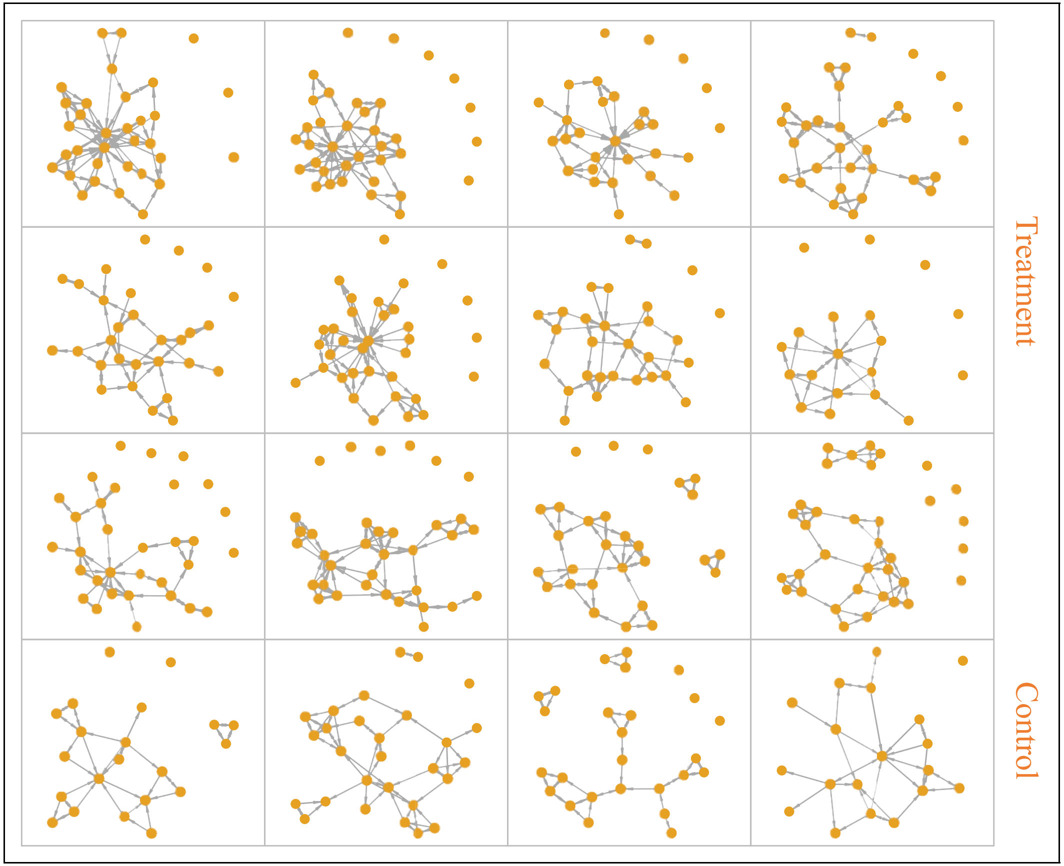

FIGURE 9. *Mathematics advice-seeking networks among teachers in eight schools the year they received coaches (treatment) and eight matched schools from the same year who did not receive coaches (control).*

policies, in which case, a nontrivial number of teachers who are not seeking advice around mathematics even when assigned a coach is not that surprising.

Thus, we define Y_k to be the change in procedural or student-centered beliefs about mathematics instruction in 1 year. X_k is an indicator for whether that school received a coach. The fitted HMMSB for Mediation is given as

$$\begin{aligned}
A_{ijk} &\sim \text{Ber}\left(S_{ijk}{}^T B R_{jik}\right)\\
S_{ijk} &\sim \text{Multi}\left(\theta_{ik}\right)\\
R_{jik} &\sim \text{Multi}\left(\theta_{jk}\right)\\
\theta_{ik} &\sim \text{Dir}\left(\xi\gamma_k\right)\\
\gamma_k &\sim N(\beta_{01} + \alpha X_k, \sigma_1^2), \gamma_k > 0\\
Y_k &\sim N(\beta_{02} + \tau' X_k + \omega\gamma_k, \sigma_2^2)\\
\beta_{01} &\sim N(0, 0.01)\\
\beta_{02} &\sim N(0, 9)\\
\alpha &\sim N(0, 9)\\
\tau' &\sim N(0, 9)\\
\omega &\sim N(0, 9)\\
\sigma_1^2 &\sim \text{Inv} - \text{Gamma}(10, 0.5)\\
\sigma_2^2 &\sim \text{Inv} - \text{Gamma}(2, 2),
\end{aligned} \tag{14}$$

where a positive Y_k value indicates that beliefs became more procedural (or student centered) during that year. As written, α is the effect of having a coach on subgroup integration and for this model, we specify B to be a 6×6 matrix with diagonal entries of 0.7 and off-diagonal entries of 0.005. Further, τ' is the effect of adding a coach on the change in beliefs conditional on the network. A positive τ' suggests that adding a coach is positively associated with beliefs becoming more procedural or more student-centered controlling for the network (subgroup insularity). Similarly, ω is the association between network structure (subgroup insularity) and the change in belief; a positive value of ω indicates that more integrated subgroups are related to beliefs becoming more procedural/ student centered.

To estimate our model given in Equation 14, we run three chains of our MCMC algorithm described in Section 5.2 and toss the first 5,000 as burn-in and retain every 500th steps for a posterior sample size of 408.

Figure 10 shows the summaries of the posterior samples for the parameters of interest. Figure 10 (top left) indicates there is a positive effect of introducing a coach on subgroup integration; schools with coaches have more integrated subgroups than schools without coaches. There is also evidence for a direct effect of

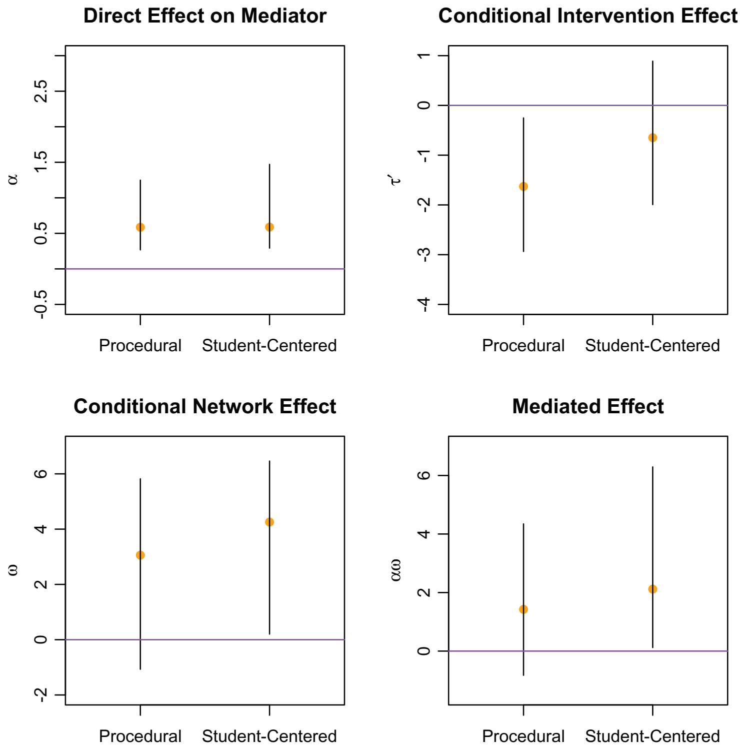

FIGURE 10. *Posterior summaries for Hierarchical Mixed Membership Stochastic Blockmodel for Mediation parameters of interest. Posterior modes and 95% equal-tailed credible intervals are given. Results suggest a positive effect of adding coaches on subgroup integration, but subgroup structure acts as a mediator on the effect of coaches on changes in student-centered beliefs only.*

adding coaches on changes in procedural beliefs but not changes in student-centered beliefs (Figure 10, top right); that is, teacher beliefs became less procedural in schools with coaches and there is not evidence that this is related to the network subgroup structure. Similarly, there is a significant association between how integrated the teacher subgroups are and change in student-centered beliefs. Teachers with more integrated subgroups became more student centered in their beliefs. Finally, Figure 10 (bottom right) shows the effect of the network as a mediator. Network subgroup insularity does appear to mediate—if we treat 95% credible intervals in the same way that we consider 95% confidence intervals—the effect of adding a coach on changes in student-centered beliefs. Network subgroup structure does not appear to significantly mediate the relationship between introducing coaches and changes in procedural beliefs.

To summarize, we found that introducing mathematics coaches affects teacher mathematics advice-seeking networks; networks with coaches generally have more integrated subgroups. This subgroup structure mediates changes in student-centered beliefs; teachers in more integrated subgroups (as a result of introducing a coach) become on average more student centered in their beliefs. Teachers with coaches also tend to become less procedural in their beliefs but there is not evidence that these changes are related to network structure.

We note that this example serves as a proof of concept but that readers should interpret our inference with caution. These data do not come from an experiment but an observational study, and although we have longitudinal data, there may be other variables to consider before concluding that coaches are the sole or even primary cause for changes in network subgroup structure or changes in beliefs.

7. Discussion

There are often studies involving multiple classrooms, schools, or other workplace organizations whose aim is to change the way people interact to ultimately change some other outcome of interest either at the individual or organization-level. In these studies, social networks operate as a natural mediator and can provide insight into the mechanisms through which these studies are effective. Motivated by this work, we proposed several ways to incorporate social networks in mediation models, so that researchers can further understand the relations between treatment and outcome.

Moreover, we extended an established statistical framework, HNMs, to include modeling networks as a mediator. We introduced the idea of incorporating networks as mediators by using parameters from a statistical social network model as the mediating variable in a standard mediation model. This generalized framework of combining social network models with mediation models will hopefully guide future methodologists in developing more models.

As a proof of concept, we also introduced the HMMSB for Mediation to both provide an example of how the HNM for Mediation framework can be used to build a model and create a model that can help researchers model interventions on networks with subgroup structure. The HMMSB for Mediation is most useful for interventions aimed at changing the level of subgroup integration in networks and those whose theory of change hypothesize that subgroup integration affects the outcome of interest.

We also provided evidence of both the feasibility and utility of the HMMSB for Mediation through a series of examples with simulated data. We found that mediation effects are recovered when the number of networks is small (20) and even when the number of nodes in each network is small (15), which suggests this model could be useful for even the small studies often observed in the social sciences. We also found that even with the addition of a rather complicated social network model, the mediation model still performs as expected.

Finally, we applied the HMMSB for Mediation to examine the effects of coaches on instructional beliefs and network structure. Practitioners generally believe that adding math coaches to a school improves the quality of mathematics instruction but the HMMSB for Mediation allows researchers to better understand the mechanisms at play. We found evidence that schools with coaches have more integrated teacher subgroups (in their advice-seeking networks) and that these integrated subgroups are also associated with teachers becoming more student centered in their beliefs about instruction. Armed with this information, researchers can better understand how instructional coaches help teachers as well as design future interventions involving coaches.

In addition, we present the HMMSB for Mediation as a first step in what we envision as a larger body of research. First, we presented a very specific model for mediation in which an intervention affects network subgroup structure. Experimental interventions may affect the ways in which individuals interact in a myriad of ways, and we argue that a large number of models can be built to explore the hypothesized changes in networks. Consider this alternative example, we might expect an intervention to shape social norms throughout a school or classroom. We therefore expect the intervention to increase the number or frequency of interactions among individuals and the increase in interaction/connections causes individuals in that system to become more similar to each other, resulting in a decrease in the variance of the outcome variable. Similarly, an intervention may impact certain individuals, such as new teachers, and we can model changes to this subset of the network using this framework.

Future work also includes the idea of extending these models to account for node-level outcomes. Even though the unit of observation (or randomization) is at the network level, we are ultimately interested in individual differences and incorporating information about the nodes and their position within their school's network as well as a node-level outcome variable is a natural next step in this work. Finally, we note that this article does not focus on causal mediation, even though mediation models are actually causal in nature. Our models have the capability of estimating causal processes, and another next step in this work is to fit these models within a potential outcomes perspective (Rubin, 2005), adapting the methods proposed by Imai, Keele, Tingley, and Yamamoto (2011) for social networks.

We conclude by situating this work into an even broader body of quantitative methodology, that is, the idea of treating a social network as the unit of observations to help address questions in the social sciences. Assuming these methods are useful to social science researchers, other analogous methods exist which also have the potential to advance social science research.

Author's Note

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the granting agencies.

Acknowledgments

I acknowledge Jim Spillane and the Distributed Leadership Studies (http://www.distributedleadership.org) at Northwestern University, including the NebraskaMATH study, for the use of their data in this work. I greatly appreciate the help provided by Jim and his colleagues.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported in part by the U.S. Department of Education, Institute of Education Sciences, under grant #R305D150045.

Notes

1. For readers interested in more rigorous statistical methods for estimating causal inference, we suggest Holland (1986) or Rubin (2005) for the potential outcomes framework and Pearl (2009) for the directed-acyclic graph perspective.
2. Note that social network ties can have any scale of measurement. We focus on binary ties, but these models apply to any type of tie.
3. As discussed in Sweet and Zheng (2017), we use γ as our measure of subgroup insularity over other measures since γ is not impacted by node membership. Measures of subgroup insularity such as the E-I index (Krackhardt & Stern, 1988) or modularity (Newman & Girvan, 2004) depend on node assignment to subgroups; in fact, one way to detect communities is to assign nodes to clusters to optimize modularity. When node membership is not clear, how these nodes are assigned impacts subgroup insularity measures; the parameter γ is independent of node assignment since nodes belong to all clusters with varying probability.
4. R code is available upon request; an R package with more efficient code is in preparation. Note that these models currently take approximately 2 hours on a standard desktop to fit.
5. Note that a single data set consists of multiple networks.
6. For more details about these data, see www.distributedleadership.org.

References

Airoldi, E., Blei, D. M., Erosheva, E., & Fienberg, S. (Eds.). (2014). *Handbook on mixed membership models and their applications*. New York, NY: CRC press.

Airoldi, E., Blei, D. M., Fienberg, S., & Xing, E. (2008). Mixed membership stochastic blockmodels. *The Journal of Machine Learning Research*, *9*, 1981–2014.

Anderson, C., & Wasserman, S. (1992). Building stochastic blockmodels. *Social Networks*, *14*, 137–161.

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.

Coburn, C. E., & Russell, J. L. (2008). District policy and teachers' social networks. *Educational Evaluation and Policy Analysis*, *30*, 203–235.

Daly, A. J., Moolenaar, N. M., Bolivar, J. M., & Burke, P. (2010). Relationships in reform: The role of teachers' social networks. *Journal of Educational Administration*, *48*, 359–391.

Erosheva, E. (2003). Bayesian estimation of the grade of membership model. *Bayesian Statistics*, *7*, 501–510.

Erosheva, E., Fienberg, S., & Lafferty, J. (2004). Mixed-membership models of scientific publications. *Proceedings of the National Academy of Sciences of the United States of America*, *101*, 5220.

Espelage, D., Holt, M., & Henkel, R. (2003). Examination of peer-group contextual effects on aggression during early adolescence. *Child Development*, *74*, 205–220.

Frank, K. A., Muller, C., & Mueller, A. S. (2013). The embeddedness of adolescent friendship nominations: The formation of social capital in emergent network structures. *American Journal of Sociology*, *119*, 216–253.

Frank, K. A., Zhao, Y., & Borman, K. (2004). Social capital and the diffusion of innovations within organizations: The case of computer technology in schools. *Sociology of Education*, *77*, 148–171.

Galyardt, A. (2012). *Mixed membership distributions with applications to modeling multiple strategy usage* (PhD thesis). Pittsburgh, PA: Carnegie Mellon University.

Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (2013). *Bayesian data analysis* (3rd ed.). Boca Raton, FL: Taylor & Francis.

Gest, S. D., & Rodkin, P. C. (2011). Teaching practices and elementary classroom peer ecologies. *Journal of Applied Developmental Psychology*, *32*, 288–296.

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.

Holland, P. W. (1986). Statistics and causal inference. *Journal of the American statistical Association*, *81*, 945–960.

Holland, P. W., Laskey, K., & Leinhardt, S. (1983). Stochastic blockmodels: First steps. *Social Networks*, *5*, 109–137.

Hopkins, M., Lowenhaupt, R., & Sweet, T. M. (2015). Organizing instruction in new immigrant destinations: District infrastructure and subject-specific school practice. *American Educational Research Journal*, *52*, 408–439.

Hopkins, M., Ozimek, D., & Sweet, T. M. (2017). Mathematics coaching and instructional reform: Individual and collective change. *The Journal of Mathematical Behavior*, *46*, 215–230.

Imai, K., Keele, L., Tingley, D., & Yamamoto, T. (2011). Unpacking the black box of causality: Learning about causal mechanisms from experimental and observational studies. *American Political Science Review*, *105*, 765–789.

Krackhardt, D., & Stern, R. (1988). Informal networks and organizational crises: An experimental simulation. *Social Psychology Quarterly*, *51*, 123–140.

Krull, J. L., & MacKinnon, D. P. (1999). Multilevel mediation modeling in group-based intervention studies. *Evaluation Review*, *23*, 418–444.

MacKinnon, D. P. (2008). *Introduction to statistical mediation analysis*. New York, NY: Routledge.

MacKinnon, D. P., & Fairchild, A. J. (2009). Current directions in mediation analysis. *Current Directions in Psychological Science*, *18*, 16–20.

Manrique-Vallier, D. (2014). Longitudinal mixed membership trajectory models for disability survey data. *Annals of Applied Statistics*, *8*, 2268–2291.

Moolenaar, N. M., Daly, A., & Sleegers, P. (2010). Occupying the principal position: Examining relationships between transformational leadership, social network position, and schools' innovative climate. *Educational Administration Quarterly*, *46*, 623.

Moolenaar, N. M., Sleegers, P. J., & Daly, A. J. (2012). Teaming up: Linking collaboration networks, collective efficacy, and student achievement. *Teaching and Teacher Education*, *28*, 251–262.

Newman, M. E., & Girvan, M. (2004). Finding and evaluating community structure in networks. *Physical Review E*, *69*, 026113.

Paluck, E. L., Shepherd, H., & Aronow, P. M. (2016). Changing climates of conflict: A social network experiment in 56 schools. *Proceedings of the National Academy of Sciences*, *113*, 566–571.

Pearl, J. (2009). Causal inference in statistics: An overview. *Statistics Surveys*, *3*, 96–146.

Penuel, W., Riel, M., Joshi, A., Pearlman, L., Kim, C., & Frank, K. (2010). The alignment of the informal and formal organizational supports for reform: Implications for improving teaching in schools. *Educational Administration Quarterly*, *46*, 57–95.

R Core Team. (2018). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.

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].

Rubin, D. B. (2005). Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, *100*, 322–331.

Snijders, T., & Kenny, D. (1999). The social relations model for family data: A multilevel approach. *Personal Relationships*, *6*, 471–486.

Spillane, J. P., & Hopkins, M. (2013). Organizing for instruction in education systems and school organizations: How the subject matters. *Journal of Curriculum Studies*, *45*, 721–747.

Spillane, J. P, Hopkins, M., & Sweet, T. M. (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., Kim, C. M., & Frank, K. A. (2012). Instructional advice and information providing and receiving behavior in elementary schools exploring tie formation as a building block in social capital development. *American Educational Research Journal*, *119*, 72–102.

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*, *90*, 149–171.

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.

Sweet, T. M., Thomas, A. C., & Junker, B. W. (2014). Hierarchical mixed membership stochastic blockmodels for multiple networks and experimental interventions. In E. Airoldi, D. Blei, E. Erosheva, & S. Fienberg (Eds.), *Handbook on mixed membership models and their applications* (Chapter 22, pp. 463–488). Boca Raton, FL: Chapman & Hall/CRC.

Sweet, T. M., & Zheng, Q. (2017). A mixed membership model-based measure for subgroup integration in social networks. *Social Networks*, *48*, 169–180.

Sweet, T. M., & Zheng, Q. (2018). Estimating the effects of network covariates on subgroup insularity with a hierarchical mixed membership stochastic blockmodel. *Social Networks*, *52*, 100–114.

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. 4024–4198). Alexandria, VA: American Statistical Association.

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 p_2 model. *Methodology: European Journal of Research Methods for the Behavioral and Social Sciences*, *2*, 42–47.

Author

TRACY M. SWEET is an assistant professor in the Department of Human Development and Quantitative Methodology, University of Maryland; email: tsweet@umd.edu. She received her PhD in statistics from Carnegie Mellon University. Her research focuses on developing statistical social network models for education applications. She is currently working on methods for modeling networks as mediators, social influence through latent variables, and methods for exploratory data analysis with multiple networks.

Manuscript received July 13, 2017
First revision received April 2, 2018
Second revision received August 24, 2018
Accepted October 15, 2018