

LATENT SPACE MODELS FOR MULTIVIEW NETWORK DATA

BY MICHAEL SALTER-TOWNSHEND AND TYLER H. MCCORMICK

University of Oxford and University of Washington

Social relationships consist of interactions along multiple dimensions. In social networks, this means that individuals form multiple types of relationships with the same person (an individual will not trust all of his/her acquaintances, for example). Statistical models for these data require understanding two related types of dependence structure: (i) structure within each relationship type, or network view, and (ii) the association between views. In this paper we propose a statistical framework that parsimoniously represents dependence between relationship types while also maintaining enough flexibility to allow individuals to serve different roles in different relationship types. Our approach builds on work on latent space models for networks (see Hoff *et al.* (2002), for example). These models represent the propensity for two individuals to form edges as conditionally independent given the distance between the individuals in an unobserved social space. Our work departs from previous work in this area by representing dependence structure between network views through a Multivariate Bernoulli likelihood, providing a representation of between-view association. This approach infers correlations between views not explained by the latent space model. Using our method, we explore 6 multiview network structures across 75 villages in rural southern Karnataka, India (Banerjee *et al.*, 2013).

1. Introduction. Understanding structure in social networks is essential to appreciating the nuances of human behavior and is an active area of research in the social sciences. Typically, human interactions occur on multiple dimensions. Individuals may be friends and co-workers, for example. The same individuals may also share membership in the same religious or professional organization. In this paper, we present a statistical model designed to glean structure from social network data collected about multiple relationships, or views. Our model builds on an active line of literature (see Hoff *et al.* (2002)) which uses low-dimensional geometric projections to represent the (likely high-dimensional) dependence structure in network data. We draw on recent and classic work on models for multivariate binary data to encode dependence between network views in

Keywords and phrases: Latent space model, Multiview relational data, Social network

the likelihood. This approach models dependence within a network view using the latent space model, while also expressing association between network views.

In particular, we address two primary statistical challenges.

1. First, our model must represent dependence structure within each view. A fundamental challenge in analyzing social network data arises because dependence between individuals' responses violates independence assumptions underlying traditional models. We wish to model the dependence that is unique to each relation.
2. Second, we wish to model associations between views. Conditional on the structure within a particular view, our model should also represent dependence between tie probabilities across views. We require this summary to be parsimonious so that we can easily compare structure across multiple similar multiview networks.

Our approach to addressing these challenges is motivated by data, presented in Section 1.1, consisting of social and financial relationships between households in 75 villages in India. Village composition differs on several key observed covariates, including religion, language and ethnicity. For a given village, we would like to understand how these covariates relate to different types of relationships. Recent work in the United States (e.g. DiPrete *et al.* (2011)), compares polarization in two network views: core association networks and broader acquaintanceship networks. DiPrete *et al.* (2011) write, for example, that “acquaintanceship networks are at least as segregated as are core networks.” The Karnataka data contains a much more diverse set of views that were available to DiPrete *et al.* (2011). It is possible to explore, for example, how homophily based on religion in social relationships compares to the level of polarization in practical relationships such as lending goods.

The Karnataka data also present a unique opportunity to explore structure between villages. Since communication between villages is negligible, these data can be conceptualized as multiple “realizations” of a population-level model of social network formation. Using this “ $n = 75$ sample” of networks, then, we would like to compare structure across villages, particularly as this structure relates to observable covariates and the relationship between views. Though it is rare to observe multiple nearly independent networks, we contend that this perspective applies more broadly to large graphs comprised of multiple, weakly connected communities.

Given the richness of these data, we propose a model that represents

structure on multiple scales. The substantive questions of interest differ depending on the scale of analysis. At the finest level, within a single view in a given village, our goal is to understand the dynamics of social structure as a function of observed household demographic differences. At the next level, we wish to understand associations across views for a given village. To truly leverage the richness of this data, we also require a measure that compares relations between views *across* villages. At this highest level we would also like to understand how associations between views differ based on observable village characteristics.

With these data in mind, we now return to the challenges that arise when modeling both within and between view structure. We also present relevant related work. First, a key statistical challenge for modeling social network data (either with multiple views or a single view) is representing the likely high dimensional dependence structure in the data in a parsimonious and interpretable way. This dependence occurs because the propensity for any two individuals to form a network relation, or edge, depends on the other edges in the network. A common network property known as transitivity, for example, implies that “a friend of a friend is likely a friend.” This feature of network data means that statistical models developed for independent observations are not appropriate. One modeling approach, the latent space model (Hoff *et al.*, 2002), represents this high-dimensional structure through a projection onto a lower dimensional *latent social space*. The latent social space, according to Hoff *et al.* (2002), represents “the space of unobserved latent characteristics that represent potential transitive tendencies in network relations.” The geometry of the social space becomes a modeling decision with substantive consequences. A latent space defined with a geometry and distance measure that satisfy the triangle inequality, for example, encodes transitivity. Since Hoff *et al.* (2002) proposed the latent space model for network data, the framework has been adapted to include model-based clustering (Handcock *et al.*, 2007; Krivitsky *et al.*, 2009) and indirectly observed network data (McCormick and Zheng, 2012). Variational approximations also facilitate using the latent space approach on larger networks (Salter-Townshend and Murphy, 2013).

Perhaps the closest alternative to the latent space approach are the various stochastic blockmodel methods. In the simplest stochastic blockmodel, each node belongs to a block or group; there are a low number of blocks and the probability of an edge between two nodes depends on the block

memberships of the nodes. Each pair of blocks has a fixed rate of edge formation between them and each block also has an internal edge formation rate. Inference then concentrates on learning the block membership vectors of each node rather than the latent positions that is our focus. The mixed-membership stochastic blockmodel (MMSB) introduced by Airoldi *et al.* (2008) extends this model so that the nodes select from a probability of block membership vector for each potential node formation; thus they exhibit membership of different blocks when interacting with different other nodes. Airoldi *et al.* (2008) demonstrate fast approximate variational Bayesian inference for the model. This model has been shown to fit well to a wide variety of real world network dataset, including document networks (Chang *et al.*, 2010) and protein-protein interactions (Airoldi *et al.*, 2006). The latent position cluster model (Handcock *et al.*, 2007) extension to the latent space model is perhaps closest of all to the stochastic blockmodel as the cluster memberships, cluster-specific spreads and inter-cluster distances correspond closely to the block memberships, internal rate of edge formation and inter-block edge formation rates respectively. Latent space models necessarily cluster together actors with connections to each other (affiliation) but the stochastic blockmodel is more flexible in that low internal edge formation rates can be chosen so that there are fewer intra-block than inter-block edges. We choose to pursue the latent space model here as we wish to enforce affiliation and because of the high transitivity exhibited by the networks in our motivating dataset, however we believe that extension of the MMSB to the multiview setting in a similar way to this paper would be both straightforward and interesting.

The second main statistical challenge in modeling multiview network data arises when considering the relationship between different types of network connections. We focus again on latent space models and review recent statistical work for multiview data in detail in Section 2. We can summarize this rapidly growing literature as two general approaches. The first set of methods accounts for multiple relations through additional structure in the latent space. This approach amounts to adding additional structure to the residual term. A second approach focuses on the data, using low dimensional representations of similarity to construct “aggregate” networks.

We present an alternative, fundamentally different, approach to modeling the relationship between network views. More formally, we model the vector of responses for each pair of individuals as arising from a distribu-

tion, which we refer to as the Multivariate Bernoulli Distribution (MVB), for multivariate binary data. We note that this model differs from a multinomial representation because individuals are allowed (and in some cases expected) to have ties across multiple views, meaning the marginal tie probabilities across all views are not restricted to sum to one (or any other total). To account for dependence between dyads, we use a conditional latent space model for each network view. This combination provides a parsimonious representation of the association between views while also encoding network structure within each view. A simple representation of association between views is essential in our case to represent patterns of associations across the 75 villages.

There is an extensive literature on statistical models for multivariate binary data. The model we present is most related to work arising from classical literature in loglinear models (see Cox (1972), for example). The likelihood framework we present was first described in this literature, then presented again recently in Dai *et al.* (2013) where the authors prove statistical properties of the MVB distribution. These models are described, among other places, in the seminal text of Bishop *et al.* (1975) and more recently in Wakefield (2013).

Previous work has also explored multivariate likelihoods for network data. Fienberg *et al.* (1985), for example, model multiview network data using a generalization of the p_1 models presented in Holland and Leinhardt (1981). Our approach differs from this earlier work in the way that we represent structure in the network, opting for the more flexible latent space representation rather than parameterizing in terms of network attributes. The p_1 model also explicitly includes both sender and receiver effects and, thus, is designed for directed networks. Our main motivating example for this work, presented in the next section, is undirected. Fienberg *et al.* (1985) also focus on data which fall in to natural, labelled subgroups, which we do not expect *a-priori* in our case. Additionally, our main focus in this work is on the relationship between network views. Despite modeling multiview structure, the Fienberg *et al.* (1985) model does not provide a simple representation for between-view dependence. To further explore the relationship between our work and the Fienberg *et al.* (1985) model, we would like to fit both models to the same data. As mentioned, this is not possible using the data which motivate our work. We instead provide a comparison using an alternate data source, the Sampson (1969) data used in the Fienberg *et al.* (1985) paper, in the Online Supplement.

More recent work (Pattison and Wasserman, 1999) includes additional network features, but compares similarly. They note that “there is a likely dependence between different ties linking any given pair of individuals. The essence of the claim is that the presence of one type of tie between individuals is likely to affect the presence of other types of tie.” They discuss various options for modeling such dependence and, like us, they employ conditional log-odds models albeit in the context of p^* or Exponential Random Graph Models (ERMGs). The key element in modeling non-independent multiview networks in Pattison and Wasserman (1999) is to model a p^* for each pair of intersecting views as well as each individual view. They condition on the complement adjacency matrix for each view at each dyad ij (that is the sociomatrix with entry ij set to undefined) and proceed using maximum pseudolikelihood estimation (MPLE; see Robins *et al.* (2007) or Strauss and Ikeda (1990) for the original work). This inference is performed jointly for all views with the change in network summary statistics calculated simultaneously for both individual views and the intersecting pairs that include that view. This approach produces a “multiplexity” coefficient that is similar to our between-view association parameter. They note that “the small frequency of $B \cap N$ ties implies that the MPLE of its corresponding parameter is likely to have a large standard error.” Further, their framework estimates the overall level of “multiplexity” in the array. In our example we wish to disentangle the relationships between specific view, which we could not accomplish with a single overall measure of association.

Another possible approach for modeling these data arises through factorization. Using recent work such as Hoff (2011a) or Hoff (2011b), we could view these data as multiway arrays. A low rank approximation to the array, then, would provide insights about social structure. This approach is again fundamentally different from our work in that we model between view dependencies as part of the likelihood, rather than as part of the latent structure. In their example with longitudinal relational data, for example, Hoff (2011a) model the relations between actors at a given time using a conditionally independent ordered probit likelihood. Instead, we opt to view relations within a view as conditionally independent, but explicitly model association structure between views in the likelihood. This distinction provides guidelines for the use of both models. Hoff (2011a) states the model components associated with multiview structure are designed to capture “heterogeneity” across views. The goal of our approach,

in contrast, is modeling “association” between views. Hoff (2011a) has a rich set of covariates about nodes (nations in that case) and multiview structure that arises from networks observed at multiple time points, a situation where ascertaining variation could be the main goal. In our case, however, we have a rich set of village-level covariates and wish to understand how those characteristics relate to associations between views.

In the remainder of this section, we present details of data collected as part of a micro finance experiment in Karnataka, a state in southern India. We are interested in a joint characterization of between network view dependencies and within network view structure. Section 2 presents recent work on latent space models for networks, providing further context for our proposed novel model, which is described in Section 3. We implement our model on the Karnataka data in Section 4. Section 5 provides a discussion and addresses future challenges.

1.1. *Multiview network data from Karnataka, India.* We examine data consisting of multiple network views collected from 75 villages in rural southern Karnataka, India. The data were collected as part of a micro finance experiment and, thus, contain both views related to social and familial interactions (being in the same family or attending temple together, for example) and views related to economic activity (lending money or borrowing rice/ kerosine, for example). All data in this example are undirected (i.e. the adjacency matrices are symmetric). Work in the economics literature has addressed the importance of the relationship between these views in outcomes such as sharing risks or exchanging favors. Jackson *et al.* (2012), for example, construct a measure of social support in each view and compare the relative level of support provided by each view within the same village. Jackson *et al.* (2012), find that their measure of support is consistently (72 out of 75 villages studied) higher for closer social relationships (such as visiting) than for “intangible” relationships (such as attending the same temple).

For each village, data consist of a census of all households in the village and detailed network information from a subset of village members. Villages range in size from about 350 members to about 1800, corresponding to approximately 75 to 350 households. Villages vary substantially in terms of wealth, religion, and language. Bharatha Swamukti Samsthe (a microfinance institution) identified these villages as places it planned to begin operations, with rollout happening during the data collection period. Data about participation in the microfinance program can be linked




FIG 1. Social networks within a village. Actors are arranged in identical order around the outside of the graphs with lines representing edges on a given relationship.

to the household census.

One of our goals is to describe the dependence structure between these various types of relationships between individuals. To illustrate the complexity of this task, Figure 1 shows the six network views we will model for one village. These views arise by collapsing the twelve views in the data into six major categories (see Section 4). Actors are arranged in the same order around all panels. The views differ in terms of both volume and structure. The graph representing family, for example, is as expected relatively sparse in this village. The panel representing social interactions, however, is quite dense. Describing structure across the views is also challenging. There are some persistent interactions, indicated most clearly in Figure 1 by the persistent horizontal lines about one-third of the way down the circles. A large portion of individuals, however, interact with one another on only some subset of these relations. Furthermore, each view has its own dependence structure with broad network properties, such as transitivity, playing different roles in each view. There is also considerable variability across villages.

2. Latent space models for network data. In this section we review a substantial literature in latent space models for networks. We begin by presenting the latent space model for a single network view. We then pay particular attention to latent space approaches to multiview network data.

2.1. *Models for a single relation.* Although our model for multiple network views can be used to extend any statistical model for single view networks, we will introduce it with the latent space suite of models. We therefore frame our discussion entirely in this setting. The latent space model for networks (Hoff *et al.*, 2002) assumes that the propensity for pairs of actors to form edges is conditionally independent given the positions of the actors in an unobserved low dimensional social space. That is for actors i and j with latent position vectors z_i and z_j and with a network density α ,

$$P(i \rightarrow j | \alpha, z_i, z_j) \propto f(z_i, z_j, \alpha).$$

If $y_{ij} = 1$ in the presence of an edge from i to j and $y_{ij} = 0$ in the absence of an edge from i to j , then we model the data using a logistic regression model with the latent space having a Euclidean distance measure;

$$\begin{aligned} P(Y^{(1)}, \dots, Y^{(R)} | \alpha, Z) &= \prod_{i,j} \left(\text{logit}^{-1}(\alpha - |z_i - z_j|) \right)^{y_{ij}} \\ &\quad \times \left(1 - \text{logit}^{-1}(\alpha - |z_i - z_j|) \right)^{1-y_{ij}}. \end{aligned}$$

Note that the likelihood is effectively a Bernoulli likelihood for each link and that these are conditionally independent given the latent space positions and density parameter.

2.2. *Models for multiple relations.* Consider now the case where $y_{ij}^{(r)} = 1$ if there is a edge from i to j on relation r , with data on $r = 1, \dots, R$ and $R > 1$. We assume that the relations are distinct, though in practice dependent since they are all realizations given one particular set of actors.

We might choose to model all views separately and sequentially (independence model), collapse all views into a single aggregate network or model all views as depending on the same underlying latent space variables. The aggregation model and the dependence model are closely related; in the aggregation model if any $y_{ij}^r = 1$ then the aggregated $y_{ij} = 1$ whereas in the dependence model all $y_{ij}^{(r)}$ are included in the likelihood. Our proposed method is one of several emerging approaches to handle multiple view networks that look to compromise between the two extremes of fully dependent and fully independent network views. We also briefly discuss two other emerging methods for modelling multiple view networks in this section, one of which also uses the latent space approach

and one that does not. We argue that our approach provides the clearest and most interpretable estimate of the inter-view dependence.

Independence model. Each relationship has its own, independent, latent space.

$$P(Y^{(1)}, \dots, Y^{(R)} | \alpha^{(R)}, Z^{(R)}) = \prod_r \prod_{i,j} \left(\text{logit}^{-1}(\alpha^{(r)} - |z_i^{(r)} - z_j^{(r)}|) \right)^{y_{ij}^{(r)}} \left(1 - \text{logit}^{-1}(\alpha^{(r)} - |z_i^{(r)} - z_j^{(r)}|) \right)^{1-y_{ij}^{(r)}}.$$

Note that inference on each view r may be performed independently.

Dependence model. Under this model the unobserved social structure is the same for all relations. That is,

$$P(Y^{(1)}, \dots, Y^{(R)} | \alpha, Z) = \prod_r \prod_{i,j} \left(\text{logit}^{-1}(\alpha - |z_i - z_j|) \right)^{y_{ij}^{(r)}} \left(1 - \text{logit}^{-1}(\alpha - |z_i - z_j|) \right)^{1-y_{ij}^{(r)}}.$$

Aggregation model. Under this model the various network views are collapsed to a single view such that if any $y_{ij}^{(r)} = 1$ then the aggregated $\tilde{y}_{ij} = 1$, otherwise $\tilde{y}_{ij} = 0$. Again, a single latent space is used to model the probability of a link.

$$P(Y^{(1)}, \dots, Y^{(R)} | \alpha, Z) = \prod_{i,j} \left(\text{logit}^{-1}(\alpha - |z_i - z_j|) \right)^{\tilde{y}_{ij}} \left(1 - \text{logit}^{-1}(\alpha - |z_i - z_j|) \right)^{1-\tilde{y}_{ij}},$$

with

$$\tilde{y}_{ij} = \begin{cases} 0 & \Leftrightarrow \sum_{r=1}^R y_{ij}^{(r)} = 0 \\ 1 & \Leftrightarrow \sum_{r=1}^R y_{ij}^{(r)} \geq 1 \end{cases}$$

Unified Graph Representation. Greene and Cunningham (2013) propose the creation of an aggregated single relation network based on the combination of the k-nearest neighbour sets for users derived from each network view. This is more sophisticated than the simple aggregation. In Greene and Cunningham (2013) the algorithm to create the aggregated network is given as:

1. For each view $j = 1$ to R , compute a similarity vector \vec{v}_{ij} between u_i and all other users present in that view, using the similarity measure provided for the view.
2. From the values in \vec{v}_{ij} , produce a rank vector of all other $(n - 1)$ users relative to u_i , denoted \vec{s}_{ij} . In cases where not all users are present in view j , missing users are assigned a rank of $(n'_j + 1)$, where n'_j is the number of users present in the view.
3. Stack all R rank vectors as columns, to form the $(n - 1) \times R$ rank matrix \mathbf{S}_i , and normalise the columns of this matrix to unit length.
4. Compute the SVD of \mathbf{S}_i^T , and extract the first left singular vector. Arrange the entries in this vector in descending order, to produce a ranking of all other $(n - 1)$ users. Select the k highest ranked users as the *neighbour set* of u_i .

Unlike our proposed method, the approach does not yield a clearly interpretable estimate of associations between the network views. This is of primary interest in our motivating problem of the Karnataka dataset.

Latent space joint model. Gollini and Murphy (2014) jointly models multiple network views by assuming that the probability of a node being connected with other nodes in each view is explained by a unique latent variable. This is the closest in spirit to the model we propose here insofar as it extends the latent space model to the multiple view setting. Unlike our proposed method, the approach does not yield a clearly interpretable estimate of associations between the network views. There are per-view network density scalar parameters α and a single latent space to model the underlying network structure. They employ a variational Bayesian algorithm to perform approximate Bayesian inference in the spirit of Salter-Townshend and Murphy (2013). The algorithm first finds estimates obtained from fitting a latent space to each network view independently. These are then used to find the joint posterior distribution of the “Latent Space Joint Model”. These results are then used to update the parameter estimates of each views Latent Space Model and this process is iterated until convergence. Note that Gollini and Murphy (2014) use the square of the Euclidean distance as it gives proportionally higher probability of a link between two nodes that are close than the usual Euclidean distance. It also requires one less approximation to be made to the log-likelihood in their Variational Bayes algorithm. We repeat the choice here for similar reasons.

Graph Correlation. Butts and Carley (2005) describe a method for measuring the correlation between multiple networks defined on the same nodes. As our interest lies in estimating the relationships between the multiple views in the Karnataka dataset, this simple calculation is close to what we want. However, as Appendix A shows, this is not the same as the correlations we model in the next section.

The difference is that the raw graph correlations of Butts and Carley (2005) are marginal whereas the correlations we ultimately report are from a joint model including a latent space model for each network view and are thus conditional upon the latent transitive structure inherent in each view. The graph correlations are averaged across these latent network topologies and are heavily influenced by them. Appendix A discusses the differences in more detail but the basic intuition is that our method and the graph correlations will be closely similar for networks with no underlying structure (i.e. Erdős-Rényi) but different when the views are transitive, etc.

3. Multivariate Bernoulli latent space model. We now describe the Multivariate Bernoulli latent space model in detail. The key feature of our model is a multivariate likelihood that represents dependence structure between views. We couple this likelihood with latent representations of social structure within each view via a latent space model. We should note that our approach is generalisable to almost any choice of probabilistic network model for the individual network views. The only constraint is that the probability of a link a view is a link-function of a linear sum. See Salter-Townshend *et al.* (2012) for a review.

We refer to the likelihood used for our model as the Multivariate Bernoulli distribution (MVB). As noted in Section 1, a representation of the MVB distribution is well-known in the loglinear models literature, though (to the authors’ knowledge) not referred to by a specific name, and arises again in recent work by Dai *et al.* (2013). The MVB extends the Bernoulli to multiple trials whose successes may be correlated. Note that unlike the Multinomial logit (and related models including correlation) for compositions, the total number of successes is not fixed. A Multinomial logit model where the number of trials is equal to the total number of links across the views for each dyad could also be performed. The total number of links follow a Binomial distribution on R trials and one could model the probability of a success in each trial as a function of a single latent space distance but we do not make use of that approach here.

Note that the likelihoods for the network models in Section 2 were all

of the univariate Bernoulli form. The probability of an edge between i and j is given by:

$$P(y_{ij}) = p^{y_{ij}} (1 - p)^{1-y_{ij}}.$$

The Bernoulli distribution is the most common model for binary variables, including binary network data. Various models for p are what distinguish the common approaches to probabilistic social network analysis. Our approach brings together these models for p with a rich literature on loglinear models for multivariate binary outcomes. Although we have chosen to focus on latent space models wherein p is modelled as the inverse logit of the Euclidean distances between points in a low dimensional latent space, the MVB extension we propose works just as well with any other latent variable model for p . See Salter-Townshend *et al.* (2012) for a review of such models, including worked out examples.

We adapt the MVB theory of Dai *et al.* (2013) to the network setting, replacing the Bernoulli likelihood with the MVB. Thus association between views enters the model through the likelihood and we will perform Bayesian inference on the association terms as well as the latent space variables. For each vector $y_{ij}^{(1)}, \dots, y_{ij}^{(R)}$, the MVB distribution explicitly allows correlation across relationships. The MVB distribution is an exponential family distribution and, thus, much of the model development remains conceptually unchanged. The model could also be generalized to continuous or valued associations using a multivariate Gaussian distribution.

We now define an MVB distribution for a single pair i, j .

Let $y_{ij} = (y_{ij}^{(1)}, y_{ij}^{(2)}, \dots, y_{ij}^{(R)})$, then

$$\begin{aligned} p(y_{ij}) &= p_{0\dots 00} (\prod_{r=1}^R (1 - y_{ij}^{(r)})) \\ &\quad \times p_{10\dots 0} (y_{ij}^{(1)} \prod_{r=2}^R (1 - y_{ij}^{(r)})) \\ &\quad \times p_{01\dots 0} (y_{ij}^{(2)} (1 - y_{ij}^{(1)}) \prod_{r \neq 2}^R (1 - y_{ij}^{(r)})) \\ &\quad \times \dots \\ (1) \quad &\quad \times p_{11\dots 1} (\prod_{r=1}^R y_{ij}^{(r)}). \end{aligned}$$

Note that for two relations, the probabilities $p_{10}, p_{00}, p_{01}, p_{11}$ define the probability that $(y^{(1)} = 1, y^{(2)} = 0)$ and so forth. The p s are constrained to sum to unity and so we perform inference on unconstrained transformations of these variables. Following Dai *et al.* (2013), we refer to these as the natural parameters. For now we will concentrate on two latent spaces and

two relations. Extension to more is straightforward. We will also assume that each latent space is responsible for each relation in turn.

The natural parameters are given by:

$$\begin{aligned} f_1 &= \log \left(\frac{p_{10}}{p_{00}} \right) \\ f_2 &= \log \left(\frac{p_{01}}{p_{00}} \right) \\ \phi_{12} &= \log \left(\frac{p_{00}p_{11}}{p_{10}p_{01}} \right) \end{aligned}$$

Examination of the terms shows that f_1 is the log odds of getting a 1 in the first relation only over getting two zeros, f_2 is the log odds of getting a 1 in the second relation only over getting two zeros and ϕ_{12} is related to the log of the covariance (see Section C). We will refer to such doubly subscripted natural parameters ϕ as association parameters. Given R views of a network we have $\binom{R}{2}$ such association parameters.

Due to the constraint that $p_{00} + p_{10} + p_{01} + p_{11} = 1$ we can also calculate all p s from these f s:

$$\begin{aligned} p_{00} &= \frac{1}{1 + \exp(f_1) + \exp(f_2) + \exp(f_1 + f_2 + \phi_{12})} \\ p_{10} &= \frac{\exp(f_1)}{1 + \exp(f_1) + \exp(f_2) + \exp(f_1 + f_2 + \phi_{12})} \\ p_{01} &= \frac{\exp(f_2)}{1 + \exp(f_1) + \exp(f_2) + \exp(f_1 + f_2 + \phi_{12})} \\ p_{11} &= \frac{\exp(f_1 + f_2 + \phi_{12})}{1 + \exp(f_1) + \exp(f_2) + \exp(f_1 + f_2 + \phi_{12})} \end{aligned}$$

In the general case of R views, the probability for all possible vectors of edges for a dyad is given by Equation 1 with each component given by:

$$(2) \quad p_{e_1 \dots e_R} = \frac{\exp(\sum_{r=1}^R (e_r f_r) + \sum_{1 \leq s < t \leq R} (e_s e_t \phi_{st}))}{1 + \sum_r \exp(f_r) + \exp(\sum_{r=1}^R f_r + \sum_{1 \leq s < t \leq R} \phi_{st})}$$

where e_r are indicator variables where a zero denotes no edge observed in view r and a 1 implies there is an edge. f_r is the log odds of a 1 in the r^{th} relation over all relations being zeros and ϕ_{st} is the association between views s and t . This equation allows us to calculate the probability of an observed edge from the set of natural parameters. We can also calculate

the natural parameters uniquely from the probability vectors although that calculation is not required for our inference.

The correlation between two views is given by

$$(3) \quad \rho = \frac{p_{00}p_{11} - p_{10}p_{01}}{\sqrt{(p_{00} + p_{10})(p_{00} + p_{01})(p_{11} + p_{10})(p_{11} + p_{01})}}.$$

The exponent of ϕ_{12} will be 1 for uncorrelated networks, less than 1 for negatively correlated networks and greater than 1 for positively correlated networks. Thus ϕ_{12} will be 0 for uncorrelated networks, negative for negatively correlated networks and positive for positively correlated networks. Appendix C demonstrates the relationship between ϕ , the network density, and ρ graphically and shows why ρ will most often take small values even when there is a strong inter-network relationship.

Our approach will be to fit a latent space model for each of the f natural parameters. The interpretation of these latent spaces is conditional on the absence of ties in all other views. In sparse graphs, these can be viewed as approximating the latent spaces for marginal tie probabilities. We discuss the importance of this nuanced interpretation further in our conclusion. We then estimate a single ϕ parameter for each pair of relations in a given network. This amounts to estimating a completely unrestricted covariance matrix between the relations. For very large numbers of relations we may wish to consider limiting consideration to specific types of covariance matrices. We also model between view associations hierarchically as a function of observed village-level covariates.

Data Generation Mechanism. We now describe the data generating mechanism for the latent space component of the model. We can sample the adjacency matrices for the views given our model as follows:

1. for $r = \{1, \dots, R\}$:
 - (a) Draw intercepts $\alpha^{(r)}$.
 - (b) for all i nodes:
 - i. Draw latent positions $z_i^{(r)}$.
2. for $1 \leq s < t \leq R$:
 - (a) Draw association parameter ϕ_{st} .
3. for all $i, j \neq i$ dyads:
 - (a) for $r = \{1, \dots, R\}$:

- i. $f_r = \alpha^{(r)} - |z_i^{(r)} - z_j^{(r)}|^2$.
- (b) Calculate all p values using Equation 2.
- (c) Simulate $\{y_{ij}^{(1)}, \dots, y_{ij}^{(R)}\}$ via a Multivariate Bernoulli distribution with parameter vector p .

Note that as per Section 2.2, we have chosen to model the log-odds of a link as linearly dependent on the squared Euclidean distances. Another approach would be to fit the latent space model to each view independently and to look at the difference between the correlations in latent distances and graph correlations. This is of course more ad-hoc than our joint modeling of the latent positions and graph associations. Further, we show empirical evidence of the superiority of our method through a cross-validation simulation study provided in the Online Supplement.

Regression Model. A main goal of our analysis is understanding how the association between views differs based on observable village-level covariates. We accomplish this using a hierarchical regression model where the outcome is the village-level association parameter and covariates are village level observable characteristics. We expect that there will be variability in the village level association parameters. To incorporate this uncertainty into our regression estimates, we fit the regression as part of a unified hierarchical model. Conditional on the village level parameters, we fit a Bayesian regression model with diffuse Gaussian priors on each coefficient. Our sampler then marginalizes over the variation in the association parameters, incorporating this variation into our resulting regression interval estimates.

Full Posterior. Bringing together the previous two components of the model, the full posterior for our model given the data is:

$$\pi(\alpha, Z, \phi | Y) \propto \pi(Y | p) \pi(\alpha) \pi(Z) \pi(\phi).$$

Each $y_{ij}^{(r)}$ is Bernoulli given the r^{th} entry of the probability vector p . For each ij dyad p is calculated from the intercepts (α), the latent positions of nodes i and j in each of the r^{th} latent space dimensions ($Z_i^{(r)}$ and $Z_j^{(r)}$) and the association parameters (ϕ) using Equation 2. The values of f are calculated as in the Data Generation Mechanism above. We place uninformative Gaussian priors for the unconstrained α, Z and ϕ .

4. Results using Karnataka data. In this section we provide results using the Karnataka India Data. These data are publicly available online at

<http://economics.mit.edu/faculty/eduflo/social>. As mentioned previously, a main challenge with these data is differentiating structure at multiple levels of analysis. In the results that follow, we contend that the substantive questions of interest also change depending on the analysis scale. In Section 4.1, we present visualizations of the latent space within a particular village broken down by household characteristics. For these finer scale analyses we are interested in understanding how structure within a village relates to these observable characteristics. Also of interest is the variability between views. Moving to Section 4.2, we examine how our village-level between view association parameter is related to village-level covariates. At this higher level of analysis, the goal is to understand relationships across villages, controlling for aspects of social structure unique to each village and view.

We implement our model using the emerging No-U-Turn sampler (Hoffman and Gelman, 2013), a variant of Hamiltonian Monte Carlo. The reasons for this choice are that it gives us good mixing due to the Hamiltonian calculations used to optimize the multidimensional proposals in the MCMC chain and there is already software available to run it (Stan Development Team, 2013). For each village, we ran four chains initialized using distinct, randomly selected starting values. Computation time for a particular village ranged from under an hour for the smallest village to several days for the largest village. Additional computational details and convergence diagnostics are presented in the Online Supplement. We could also have used a Variational Bayes approach such as that used in Salter-Townshend and Murphy (2013) but we wish to retain as much of the posterior dependency structure as possible.

4.1. Village level results. In this section we present a small subset of the latent space results within each village. We provide results for all villages in the Online Supplement. We fit our model to multiview data describing 6 social relations between households in Karnataka, India. As discussed previously, these data arise from 75 villages with little interaction between villages and thus provide a unique opportunity to explore how diversity in network structure across villages is related to association across views. In addition to information on network structure for multiple views, we also have household and individual level covariates that we aggregate to construct village-level characteristics. Within each village, our goal is to relate social structure for a particular view to a set of household-level covariates. Our proposed model accomplishes this through the latent space

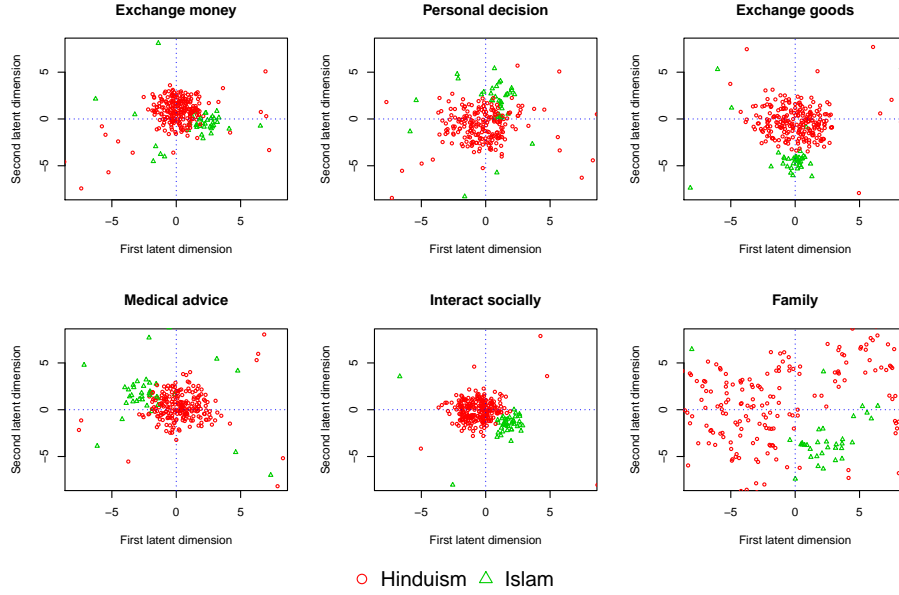


FIG 2. Latent space for Village 71 by religion. Maximum *a-posteriori* latent position for each relation, adjusted to a common orientation based on “Interact socially.”

model within each village and each view. Next, we also wish to compare the association between views across villages. The scale of this question requires a more parsimonious representation of the relationship between views.

We find that each network view in the Karnataka dataset is strongly transitive. We used `cug.test` in the R package `sna` (Butts, 2010) to determine the probability that a random graph with the same number of edges exhibits as high or higher a transitivity measure as each of the Karnataka network views. In all cases the empirical *p*-value returned was within machine precision of zero, showing that the transitivity of the graphs of each network view were not from a random graph with no structure. As per Butts (2008), the random graphs were conditioned on having the same number of edges. This result motivates our choice of a Euclidean latent space. Figures 2 and 3 show the maximum *a-posteriori* latent position for respondents in two villages. Figure 2 is colored by religion, whereas Figure 3 is colored by caste. In both figures we used a Procrustes transformation (see Hoff *et al.* (2002)) to rotate the points to an orientation that most

matches the “interact socially” view. As noted in the model description, these latent spaces are on the natural parameters of the MVB distribution. As with other latent space approaches, we must choose the dimension of the latent space beforehand. We used BIC (see supplement to paper) to evaluate the goodness-of-fit of the model using latent spaces ranging from one to five dimensions. Latent spaces of two and three dimensions yielded the best performance, with performance in the three dimensional case being slightly superior; we opted to present results with two dimensions for ease of visualization. We present these results in full detail in the Online Supplement, where we also present regression results in three dimensions to demonstrate that our substantive conclusions are qualitatively similar. If the primary goal were, for example, link prediction, we could select a latent dimension based on hold-out data experiments.

Several striking patterns emerge in the figures. First, the figures demonstrate that there is substantial variability in the latent structure across views. For these villages, the views representing social interactions and money exchange are most concentrated around the origin, giving the smallest latent distances between individuals. These smaller latent distances imply that the model is less reliant on higher order terms in estimating interaction frequency. The view encoding being related has the most diffuse latent structure. Such structure occurs because individuals linked by familiar ties form relatively small, dense clusters and there is a high propensity for ties within a cluster but few opportunities for ties between. Second, we see substantial homophily based on religion across all views in Figure 2. There is also evidence of homophily based on caste though it is more concentrated in some views than others. The interact socially and exchange money views, for example, demonstrate the most homophily based on caste in Figure 3. The view encoding familial ties displays pronounced homophily, but again remains the most diffuse of the views as there are comparatively few connections across family units even within the same caste. Third, despite applying a Procrustes transformation, the orientation of the latent spaces still varies across views. In Figure 2, for example, we see that the latent positions of individuals practicing Islam are generally to the upper right of individuals practicing Hinduism on the view related to financial advice, whereas they are to the upper left for the view measuring medical advice. With heterogeneous network structure, the Procrustes transformation is unable to match these orientations. Isolates also play a substantial role; making this observation a byproduct of sparsity in the

underlying graphs.

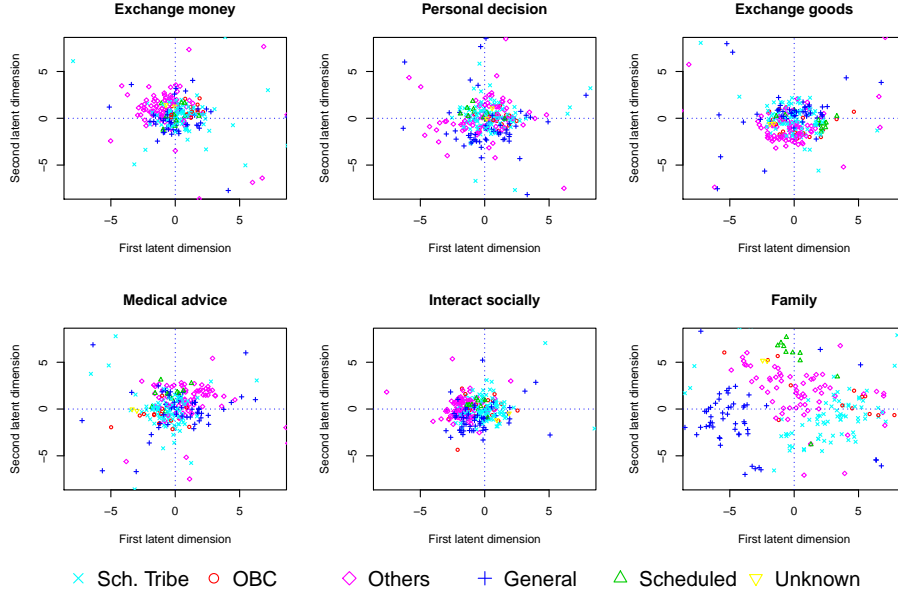


FIG 3. Latent space for Village 29 by caste. Maximum a-posteriori latent position for each relation, adjusted to a common orientation based on “Interact socially.”

Moving to the next level of analysis, we can also explore the association between views within a given village. For interpretability, we express these results in terms of the correlation parameter from the MVB. To get this parameter, we first convert the association parameter, ϕ in our model to correlation. We calculate the pairwise probability of each possible combination of links for two network views using Equation 2. The correlation between the two network views is then calculable using Equation 3. Figure 4 shows results for the raw graph correlation (Butts and Carley, 2005) and the correlation posterior means and standard deviations from our model for the 6 network views. This figure uses the same village as in Figure 2.

The proposed measure represents the excess correlation after “factoring out” social structure encoded within each network view. We contrast this conditional correlation measure with the marginal measure obtained using graph correlation. In Figure 4 we see that controlling for structure within each view can substantially change our interpretation of the association

between views. The exchange goods/interact socially pair, for example, is one of the least correlated when considering the marginal measure of graph correlation, but has the strongest association when using the proposed conditional measure. When interpreting the results, the conditional approach provides a measure that accounts for both changes in density between views and differences in higher order interaction terms. The graph correlation measure adjusts for density, but the relation between views is still confounded with network structure beyond that present in an Erdős-Rényi graph. Overall, the conditional correlation approach also produces a more pronounced distribution of associations, with few associations standing out as clearly the largest. We investigate these differences further with a simulation study in Appendix A.

The proposed metric produces correlations that are very small for the village presented in Figure 4. The absolute value of the coefficients is misleading, however, as the range of possible values depends on the structure of the graph. For both the graph correlation and our conditional approach, the overall density of the network is influential. Large negative correlation is only possible in very dense networks and large positive correlation only in very sparse networks. The presence of latent structure shifts the distribution of possible correlation values for our method to the left. Graph correlation, in contrast, does not account for higher order network structure and thus is unaffected. We explore this dependence further in Appendix C.




FIG 4. Plot of correlation (ρ) matrices for village 72 using *gcor* (Butts, 2010) (left), the posterior mean (middle) and posterior standard deviation (right) from our model. Red implies a stronger correlation but note the difference in scale. The difference under the two approaches is due to the incorporation of the latent space component in the second approach and is discussed in Appendix A. Figure C shows that large negative correlation is only possible in very dense networks and large positive correlation only in very sparse networks.




FIG 5. Association results using Karnataka data. Each plot represents a regression model for a particular view pair. Solid dots represent coefficient estimates in a model where covariates are village-level measures of socioeconomic and demographic characteristics for villages. The outcome is either the graph correlation (blue) or the association parameter from the MVB latent model (orange). The results indicate that accounting for network structure through the MVB cluster model can have a substantial impact. All variables are standardized for comparison across models. Additional results are presented in Appendix B.

4.2. *Cross-village results.* We now move to our results across villages. In contrast to the previous results where our substantive questions were focused mostly on understanding properties of local graph structure, we are now interested in broad trends in association between views based on observable village level characteristics. For each village we now have an estimated association based on the MVB latent model and a series of village-

level covariates. Our hierarchical model facilitates including a regression component on the between-view associations. That is, we construct regression models where the response variable is the MVB latent model association parameter and the predictor variables are our observed covariates. We used covariates related to the socioeconomic status (e.g. proportion of households with a latrine, proportion with a number of different types of roof, or average number of rooms per home), household characteristics (proportion of households that report having a leader), and demographics (average time in the village). We also experimented with models including other village-level covariates (including average age, language, and religion) though these models did not perform as well using measure of model fit such as AIC. For comparison, we performed the regressions using graph correlation as a response variable.

Figure 5 presents the results from our regression models for four view pairs. Since we used six network views, there are a total of fifteen view pairs. The remaining views are presented in Appendix B. For each regression coefficient, we can interpret the result using graph correlation as an increased (or decreased) propensity for people in villages with high levels of the covariate with an interaction on one view to also have an interaction on the second view. For example, the coefficients for proportion with a tile and stone roof in the pair medical advice/family indicate that in villages where more individuals have tile roofs there is an increased propensity for those who are related also share medical advice. The roof type variables can be interpreted as a crude measure of socioeconomic status with the thatched roof indicating the most impoverished individuals and RCC (concrete) among the most affluent. The thatched roof category does not follow the general trend between roof type and model coefficients. The number of households in the thatched roof category is very small (less than 2% of the data), however, so there is very little data to estimate these coefficients. On both the family/medical advice and family/exchange money associations there is a consistent trend based on roof type. Looking first at the family/money association, households in villages with (roughly speaking) lower socioeconomic status were less likely to exchange money with individuals they are related to. This trend is reversed for medical advice, however. Individuals in villages with lower status, as measured by the fraction of lower quality roofs, were more likely to share medical advice with family members. In both cases, these effects are after accounting for the nuances of within-network structure in each network.

The results using graph correlation do not control for network structure within each view, meaning that the effect of association between views is potentially confounded with local graph structure. Using the MVB latent space model, in contrast, we can explore association between views while controlling for network structure within each view. In terms of interpretation, the coefficients for the MVB latent space model now reflect the increased (or decreased) propensity for individuals in villages with a high value of a given covariate to interact on a second view once they already interact on one view in the pair, while controlling for the role individuals play within a given network. Turning back to Figure 5, we see that this distinction has significant implications in some circumstances. In the exchange money/family pair, for example, we see that the sign of the coefficients changes from positive using graph correlation to negative with the MVB latent space model. This implies that the positive association between these views for villages with large proportions of tile or stone roofs (individual with relatively low socioeconomic status) seen by graph correlation is confounded with network structure within each view. In other words, the MVB latent space results indicate that, controlling for whether individuals are related, individuals within villages with relatively low SES were less likely to also exchange money. This result is not, in contrast, due to confounding with network structure when looking at the medical advice/family pair. As expected, we also see that controlling for structure through the MVB latent space model also serves to decrease the size of coefficients in several other cases as can be seen with coefficients in the medical advice/social pair or the financial advice/social pair.

5. Discussion and conclusion. We present a statistical model for multiview network data. Our model builds on previous work on latent space models for networks as well as a deep literature in modeling discrete multivariate data. The proposed method uses a multivariate likelihood to model associations between views, which provides a parsimonious representation of between-view structure. Conditional latent space models account for dependence arising from social network structure. Importantly, ours is the only model we know of that can not only capture but also estimate negative correlation between multiple network views jointly with complex network topography. In our results we present a comprehensive comparison across all fifteen possible combinations of six views. This approach leads to testing a very large number of hypotheses, which should be considered when using this method.

The importance of multiview network models also rises with increasing ability to measure various types of networks using technology. Social media data such as Twitter or Facebook, for example, provide access to a wealth of data consisting of multiple types of communications. In Twitter, for example, users can also communicate in either passive or active relationships. “Following” in Twitter terminology allows a users to see updates from another user but typically does not require approval or interaction with the user being followed. Interactions can also be more active and conversational, however, with the individuals using the @reply command or contributing to broader conversations by including a #hashtag identifier in their tweet. Twitter users project their thoughts toward an imagined audience of networked individuals, some of whom bear reciprocal edges to the users themselves and some of whom do not. This interesting mix of public and private attention requires users to maintain a balance between transparency and authenticity in the material they choose to tweet.

The conditional latent space approach provides flexibility in modeling social network structure within each view. Our method does, however, share difficulties related to model selection that are common in latent space models. Selecting the dimensionality of the latent space remains a topic of research with these models. In our case, we use the latent space to account for within-view structure but do not emphasize interpretation of the latent representation. We fixed the dimensionality of the latent space to be the same across all views to encourage comparable interpretation across views. We experimented with multiple latent dimensions and found that the association measures were relatively insensitive, though more exploration of this area could be done for different modeling objectives. In our case we value the interpretation of the latent space approach, though we could also approach the within-view modeling using other social network models. The behavior of the MVB framework under other models and model misspecification is an open area for future work.

The conditional representation also makes interpretation of the latent distances a challenge. This issue is well-known in the loglinear models literature. Alternative models for multivariate binary data have been proposed, but these models also have substantial issues with interpretation, often in ways that are less straightforward to interpret than the likelihood used here (see Wakefield (2013) for a review). In our application, the main focus is on understanding the association between village-level covariates and views. In other applications, however, interpreting marginal tie prob-

abilities within a view may be a priority. In such instances, we can obtain posterior distributions over these probabilities during sampling, or decomposition methods could be considered (e.g. Hoff (2011a)).

6. Acknowledgements. McCormick’s work is supported by U.S. Army Research Office 62389-CS-YIP. Salter-Townshend is currently supported by NIH grant R01 HG006399. We would like to thank two anonymous Referees, the Associate Editor, and the Editor for very helpful comments. We also thank Abhijit Banerjee, Arun G. Chandrasekhar, Esther Duflo and Matthew Jackson for providing these rich data through their Social Networks and Microfinance Project.

APPENDIX A: DIFFERENCE BETWEEN MARGINAL AND CONDITIONAL GRAPH CORRELATIONS

In order to demonstrate the difference between the conditional correlations inferred by our model and raw (marginal) graph correlations, we present results on some simulated datasets. We will show that in the absence of any underlying latent structure to each network view (effectively Erdős-Rényi graphs), our method returns values close to the graph correlations. However, when there is a non flat network topography our method is preferable to such marginal correlations. We simulated 3 networks with a ground truth association matrix given by

$$\phi_{st} = \begin{bmatrix} NA & 4.42 & 3.49 \\ 4.42 & NA & -1.50 \\ 3.49 & -1.50 & NA \end{bmatrix}.$$

As this is a symmetric matrix with just 3 unique terms we will report it and related matrices as vectors of length 3 with terms corresponding to view pairs (1,2), (1,3) and (2,3) respectively. Thus the association values used to generate the data are

$$(\phi_{12}, \phi_{13}, \phi_{23}) = (4.42, 3.49, -1.5).$$

We explore two versions of the simulated data here. The first version has all latent positions in each network view set to zero (i.e. no latent space structure) and the second version has randomly generated latent positions. We report results for multiple (20) runs of both versions, each involving the generation of random data and subsequent application of gcor (Butts, 2010) and our method to infer inter-network-view correlations.

In both versions, we simulated the joint 3-network-view links as per our Multivariate Bernoulli Distribution.

A.1. With No Latent Structure. For the version with no latent structure the ground truth correlations ρ between each pair of views is constant across simulation runs and is calculated using Equation 3 to be

$$\text{ground truth: } \rho = (\rho_{12}, \rho_{13}, \rho_{23}) = (0.383, 0.349, -0.067).$$

The raw graph correlations mean and standard deviations across the 20 simulations were found to be:

$$\text{gcor: } \hat{\rho} = (0.515, 0.307, -0.001) \quad s_{\rho} = (0.027, 0.020, 0.018).$$

We then fit our model using our MCMC algorithm and obtain draws from the posterior for the latent positions Z , intercepts α and associations ϕ . We find that the mean across 20 simulations of the posterior means and standard deviations for the posterior means for the associations are

$$\text{MLSM: } \hat{\phi} = (4.661, 3.635, -1.589) \quad s_{\phi} = (0.224, 0.187, 0.160).$$

We then calculate the pairwise correlations for given values of the intercepts α and distances between latent positions Z . We average these values across the MCMC iterations Equation 3 to get:

$$\text{MLSM: } \hat{\rho} = (0.367, 0.333, -0.067) \quad s_{\rho} = (0.012, 0.013, 0.008).$$

Thus the algorithm was able to capture the correct associations and correlations across simulations.

Although the results from gcor are comparable to the results using our method for this simulated data without latent structure in the networks, we note that gcor failed to capture the correct ρ values within the range of values across the 20 simulation runs, whereas our results for posterior means span the true values for both ϕ and ρ and are centered on them. This is hardly surprising of course as we are fitting the same model used to generate the data (apart from the inclusion in the inference of latent positions which were set to zero to generate the data).

We can also examine 95% credible intervals based on the MCMC output for each of the 20 simulations. We found that these intervals included the correct values of ϕ , averaged across the 3 view pairings, 90% of the time. For ρ , it was 85%. Thus, even within a simulation, our posterior does a good job of inferring the correct association values.

A.2. With Latent Structure. We now report results for a simulation study including randomly generated latent positions associated with each network view. We observe that the ρ values found by our method now differ more substantially from the graph correlations. Note that due to the simulation of random latent space positions, the true ρ values now vary across simulation runs. The ϕ values are the same as per the simulations with no latent structure. The true (i.e. values used to generate the simulated datasets) mean and standard deviations of ρ were:

$$\text{ground truth: } \hat{\rho} = (0.052, 0.037, -0.003) \quad s_{\rho} = (0.006, 0.005, 0.000).$$

Note that the standard deviation of these values across simulations is low.

The raw graph correlations mean and standard deviations were calculated to be:

$$\text{gcor: } \hat{\rho} = (0.634, 0.439, 0.145) \quad s_{\rho} = (0.057, 0.065, 0.068).$$

These values are far removed from the ground truth, with all 3 far too strongly positive. This is because correlation due to the latent structure is dominating the correlations.

Conversely, our method found values of

$$\text{MLSM: } \hat{\rho} = (0.029, 0.018, -0.002) \quad s_{\rho} = (0.006, 0.004, 0.000),$$

which are more in line with the ground truth. Across the 20 simulation runs, we captured the ground truth value of ρ in the 95% credible interval 60% of the time. Thus, in the presence of latent structure underlying the multiple networks, the graph correlation does a far poorer job of capturing the true correlation between the views. For real datasets, such as our motivating example, we therefore advocate using a model that jointly estimates the correlations and the latent structures of the multiple views.

APPENDIX B: ADDITIONAL RESULTS FOR KARNATAKA

In this section, we present results for the remaining views not presented in Figure 5. Figure B presents plots giving point estimates and error bars for regression models where the outcome is either the graph correlation or the association parameter of the MVB latent model. Covariates are village level measures of demographic and socioeconomic composition. The results support the discussion in Section 4.

APPENDIX C: RELATIONSHIP OF CORRELATION ρ TO DENSITY α AND ASSOCIATION ϕ

It is interesting to observe how ρ varies with network density α and association ϕ . Figure C presents a contour plot of the correlation ρ_{12} against α and ϕ_{12} . Note that the distances in the latent spaces are set to zero before using Equations 2 and 3. The effect of having equal but non-zero distances is to shift the plot to the right of zero on the x-axis. If the distances are unequal in the two views, the contours also move away from zero in the y direction.

It can be seen from Figure C that the absolute correlation is largest for either high density networks with large negative association parameter or low density networks with large positive association parameter. The sign of the correlation is the same as the sign for the free association parameter but the actual correlation values will be small in magnitude even if the free parameter is large for many networks.

REFERENCES

- E. M. Airoldi, D. M. Blei, S. E. Fienberg, E. P. Xing, and T. Jaakkola. Mixed membership stochastic block models for relational data with application to protein-protein interactions. In *Proceedings of the international biometrics society annual meeting*, page I5, 2006.
- E. Airoldi, D. Blei, S. Fienberg, and E. Xing. Mixed-membership stochastic blockmodels. *Journal of Machine Learning Research*, 9:1981–2014, 2008.
- A. Banerjee, A. Chandrasekhar, E. Duflo, and M. O. Jackson. The Diffusion of Microfinance. The Abdul Latif Jameel Poverty Action Lab Dataverse, 2013.
- Y. Bishop, S. Fienberg, and P. W. Holland. *Discrete Multivariate Analysis: Theory and Practices*. MIT Press, 1975.
- C. T. Butts and K. M. Carley. Some simple algorithms for structural comparison. *Computational & Mathematical Organization Theory*, 11(4):291–305, 2005.
- C. T. Butts. Social network analysis: A methodological introduction. *Asian Journal of Social Psychology*, 11(1):13–41, 2008.
- C. T. Butts. *sna: Tools for Social Network Analysis*. University of California, Irvine, 2010. R package version 2.1-0.
- J. Chang, D. M. Blei, et al. Hierarchical relational models for document networks. *The Annals of Applied Statistics*, 4(1):124–150, 2010.
- D. R. Cox. Regression models and life-tables. *Journal of the Royal Statistical Society, Series B*, pages 187–220, 1972.
- B. Dai, S. Ding, and G. Wahba. Multivariate Bernoulli distribution. *Bernoulli*, 19(4):1465–1483, 2013.
- T. A. DiPrete, A. Gelman, T. McCormick, J. Teitler, and T. Zheng. Segregation in social networks based on acquaintanceship and trust1. *American Journal of Sociology*, 116(4):1234–83, 2011.
- S. E. Fienberg, M. M. Meyer, and S. S. Wasserman. Statistical Analysis of Multiple Sociometric Relations. *Journal of the American Statistical Association*, 80(389):51–67, 1985.

- I. Gollini and T. B. Murphy. Joint modelling of multiple network views. *Journal of Computational and Graphical Statistics*, -(to appear):00–00, 2014.
- D. Greene and P. Cunningham. Producing a unified graph representation from multiple social network views. *Proceedings of the 5th Annual ACM Web Science Conference (WebSci’13)*, pages 118–121, 2013.
- M. Handcock, A. E. Raftery, and J. M. Tantrum. Model-based clustering for social networks. *Journal of the Royal Statistical Society Series A*, 170:301–354, 2007.
- P. D. Hoff, A. E. Raftery, and M. S. Handcock. Latent space approaches to social network analysis. *Journal of the American Statistical Association*, 97:1090–1098, 2002.
- P. Hoff. Hierarchical multilinear models for multiway data. *Computational Statistics & Data Analysis*, 55:530–543, 2011.
- P. Hoff. Separable covariance arrays via the Tucker product, with applications to multivariate relational data. *Bayesian Analysis*, 6:179–196, 2011.
- M. D. Hoffman and A. Gelman. The no-U-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, in press, 2013.
- P. W. Holland and S. Leinhardt. An Exponential Family of Probability Distributions for Directed Graphs. *Journal of the American Statistical Association*, 76(373):33–50, 1981.
- M. O. Jackson, T. Rodriguez-Barraquer, and X. Tan. Social Capital and Social Quilts: Network Patterns of Favor Exchange. *American Economic Review*, 102(5):1857–1897, August 2012.
- P. N. Krivitsky, M. S. Handcock, A. E. Raftery, and P. D. Hoff. Representing degree distributions, clustering, and homophily in social networks with latent cluster random effects models. *Social Networks*, 31(3):204–213, 2009.
- T. H. McCormick and T. Zheng. Latent demographic profile estimation in hard-to-reach groups. *Annals of Applied Statistics*, 6(4):1795–1813, 2012.
- P. Pattison and S. Wasserman. Logit models and logistic regressions for social networks: II. Multivariate relations. *British Journal of Mathematical and Statistical Psychology*, 52(2):169–193, 1999.
- G. Robins, P. Pattison, Y. Kalish, and D. Lusher. An introduction to exponential random graph (p^*) models for social networks. *Social networks*, 29(2):173–191, 2007.
- M. Salter-Townshend and T. B. Murphy. Variational Bayesian inference for the latent position cluster model for network data. *Computational Statistics & Data Analysis*, 57(1):661 – 671, 2013.
- M. Salter-Townshend, A. White, I. Gollini, and T. B. Murphy. Review of statistical network analysis: models, algorithms, and software. *Statistical Analysis and Data Mining*, 5(4):243–264, 2012.
- S. F. Sampson. Crisis in a cloister. *Unpublished doctoral dissertation, Cornell University*, 1969.
- Stan Development Team. Stan: A C++ library for probability and sampling, version 2.0, 2013.
- D. Strauss and M. Ikeda. Pseudolikelihood estimation for social networks. *Journal of the American Statistical Association*, 85(409):204–212, 1990.
- J. Wakefield. *Bayesian and Frequentist Regression Methods*. Springer, New York, 2013.

E-MAIL: salter@stats.ox.ac.uk

DEPARTMENT OF STATISTICS
UNIVERSITY OF OXFORD

E-MAIL: tylermc@uw.edu

DEPARTMENT OF STATISTICS
DEPARTMENT OF SOCIOLOGY
UNIVERSITY OF WASHINGTON

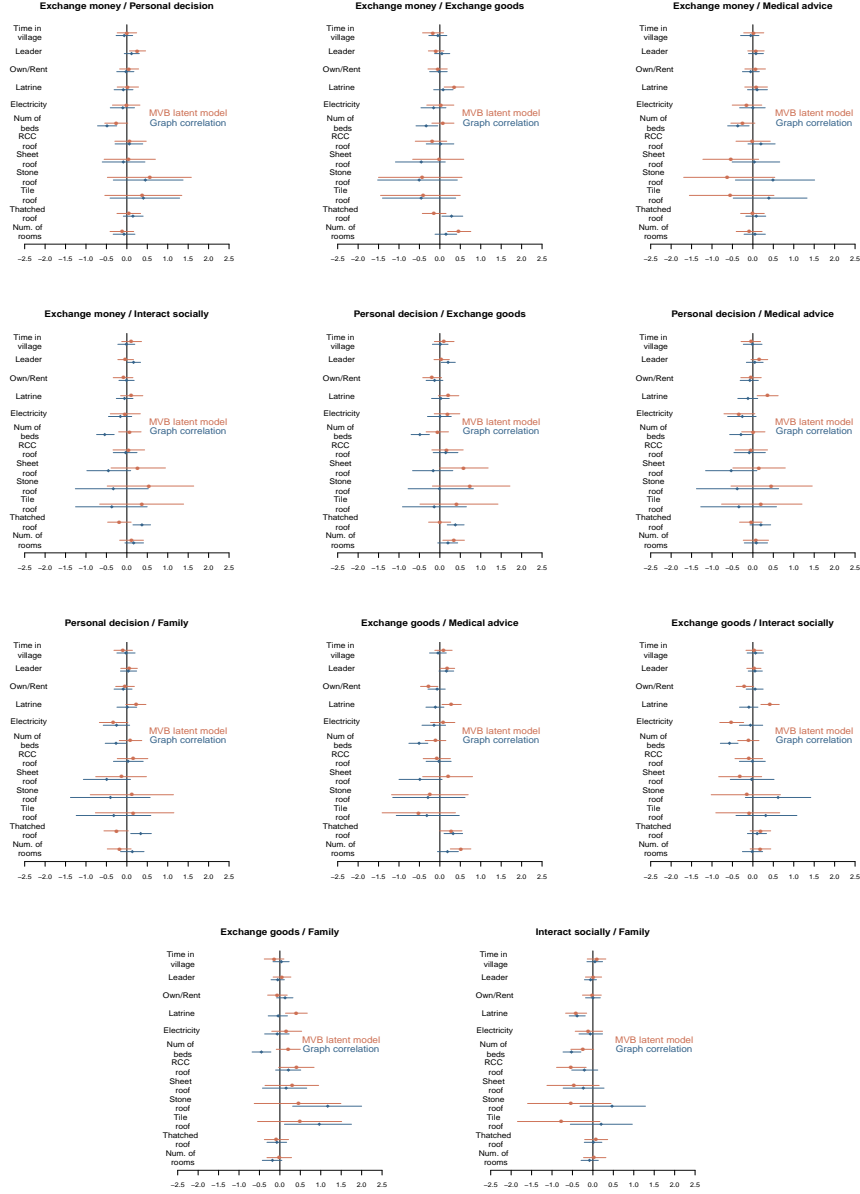


FIG B. Results for additional view pairs. These plots present coefficients and error bars for the remaining relations not presented in Figure 5. Each plot represents a single view pair. Dots represent point estimates in a regression model where the outcome is either graph correlation or the association parameter in our MVB latent model for a particular village and village level covariates. Bars represent 80% uncertainty intervals. We standardized all variables for comparison across outcomes.




FIG C. Contour plot of how the correlation ρ between two views varies as a function of the density of the networks α and the association term ϕ_{12} . Note that for sparse networks (towards the left), positive associations can lead to large positive multiview correlations; however negative correlations are limited to be small in magnitude.