A LATENT VARIABLE APPROACH FOR MODELING RELATIONAL DATA WITH MULTIPLE RECEIVERS

BY JORIS MULDER[1,a] AND PETER D. HOFF[2,b]

[1]*Department of Methodology and Statistics, Tilburg University,* [a]*j.mulder3@tilburguniversity.edu*
[2]*Department of Statistical Science, Duke University,* [b]*peter.hoff@duke.edu*

Directional relational event data, such as email data, often contain unicast messages (i.e., messages of one sender toward one receiver) and multicast messages (i.e., messages of one sender toward multiple receivers). The Enron email data that is the focus in this paper consists of 31% multicast messages. Multicast messages contain important information about the roles of actors in the network, which is needed for better understanding social interaction dynamics. In this paper a multiplicative latent factor model is proposed to analyze such relational data. For a given message, all potential receiver actors are placed on a suitability scale, and the actors are included in the receiver set whose suitability score exceeds a threshold value. Unobserved heterogeneity in the social interaction behavior is captured using a multiplicative latent factor structure with latent variables for actors (which differ for actors as senders and receivers) and latent variables for individual messages. The model is referred to as the multicast additive and multiplicative effects network (mc-amen) model. A Bayesian computational algorithm, which relies on Gibbs sampling, is proposed for model fitting. Model assessment is done using posterior predictive checks. Numerical simulations show that the model is widely applicable for various scenarios involving multicast messages. Furthermore, a mc-amen model with a two-dimensional latent variable can accurately capture the empirical distribution of the cardinality of the receiver set and the composition of the receiver sets for commonly observed messages in the Enron email data. In the Enron network, actors have a comparable (but not identical) role as a sender and as a receiver in the network.

1. Introduction. Social behavior between individuals in a network can often be characterized by short communication messages of one actor toward another actor or several other actors. In an information network of employees in an organization, a relational event can be an email message sent by an employee to one or several fellow employees (e.g., Mulder and Leenders (2019), Perry and Wolfe (2013)), in a classroom a relational event can be a teacher hushing one or several students (DuBois et al. (2013)), and in a military convoy a relational event can be a soldier giving a command to one or more fellow soldiers (Leenders, Contractor and DeChurch (2016)). Such relational data are becoming increasingly available due to the technical innovations, e.g., email, instant messaging apps, or radio communication, which have resulted in further development of statistical models for time-stamped social interaction data (e.g., Arena, Mulder and Leenders (2022), Brandes, Lerner and Snijders (2009), Butts (2008), Meijerink-Bosman et al. (2022), Meijerink-Bosman, Leenders and Mulder (2022), Quintane et al. (2014), Stadtfeld and Block (2017)).

A key property of these relational data is that they contain unicast messages (i.e., messages sent from one actor to another actor) as well as multicast messages (i.e., messages sent from one actor to multiple other actors). The Enron email data (Cohen (2009)), as compiled by Zhou et al. (2007) with $M = 21{,}635$ messages among $N = 156$ actors, which is the focus in this paper, consists of 69% unicast messages and 31% multicast messages (Supplementary
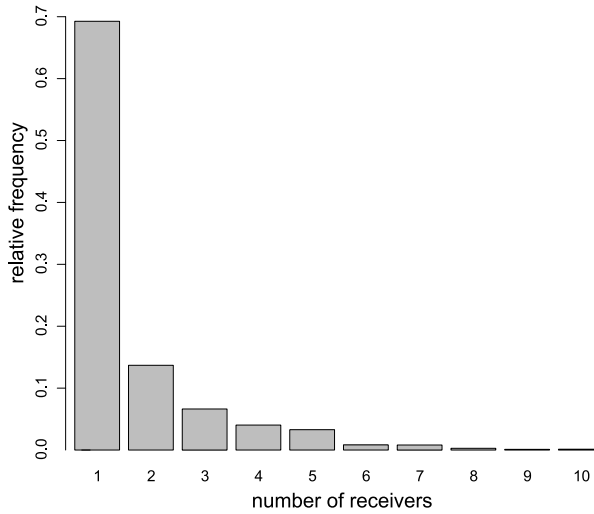
FIG. 1.   *Empirical distribution of the number of receivers per message in the Enron email data (*Cohen* (2009),
Zhou et al. (2007)).*

Material). The empirical distribution of the cardinality of the receiver set over the entire data
is presented in Figure 1. The distribution of the number of receivers has a typical shape that is
often observed with email data in practice (e.g., see also the email data from Eckmann, Moses
and Sergi (2004) and Mulder and Leenders (2019)) where unicast messages are observed
most frequently, and the observed number of multicast messages tends to decrease as the size
of the receiver set grows.

Even though multicast relational data have not received a lot of attention in the statistical
literature, multicast messages reveal important insights about social relations between actors,
about social hierarchy of actors, and about the unique role of actors in a network. To make
this more concrete, we consider a subset of the messages sent between actors 26, 49, 53, 59,
and 130 in the Enron email network (all from the legal department) who sent 24, 77, 13, 408,
and 102 messages among each other, respectively, out of which 70.5% are unicast messages
and 29.5% multicast messages (similar as in the full data set). The observed frequencies of
the receiver sets of the most active sender (actor 59), with 408 sent messages, can be found
in Figure 2. First, we see that actor 130 is mainly popular as a receiver in unicast messages
and this actor is practically absent in multicast messages. This suggests that actor 130 has
a unique relation with this sender which is different from the other actors. It may be that
actor 59 and actor 130 are both involved in a specific project or topic which is not relevant
for the other actors. Second, we see that actors 49 and 53 have a similar role in terms of their
relative popularity in unicast messages (where actor 49 is slightly more popular) as well as in
multicast messages (where they generally appear both in frequently observed receiver sets,
namely, {49, 53} and {26, 49, 53}). This suggests that actors 49 and 53 have a similar role or
hierarchical level in the network in relation to actor 59. Third, we can see that actor 26 has a
somewhat subordinate role (relative to actor 49 and actor 53), as actor 26 is mainly popular
as a receiver when actors 49 and 53 are also in the receiver set, but actor 26 is not a popular
receiver in unicast messages.

Because most statistical approaches for modeling relational data are designed for dyadic
observations involving two actors, one could split all multicast messages into multiple unicast
messages of the sender to the different receivers. Effectively, we would then be left with the
frequencies of how often each actor appears as a receiver for every given sender and thereby
loose the unique information about the specific roles of actors contained in multicast mes-
sages. For example, for the above mentioned subset of messages that were sent by actor 59,
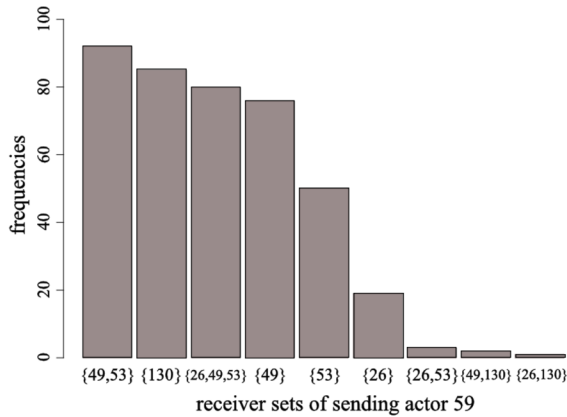
FIG. 2. *Observed frequencies of the receiver sets of messages sent by actor* 59 *in a subset of the Enron data consisting of actors* 26, 49, 53, 59, *and* 130.

this separation step of the multicast messages would have resulted in observed inclusion proportions of 0.25, 0.61, 0.55, and 0.22 for actors 26, 49, 53, and 130, respectively. Based on these proportions, we can only conclude that actors 26 and 130 are less popular receivers than actors 49 and 53. We no longer see the unique role of actor 130, the similar roles of actor 49 and 53, or the subordinate role of actor 26 relative to actors 49 and 53. Another issue of the separation step is that no information is available any more about the empirical distribution of the number of receivers in the data (such as in Figure 1), making it problematic (or practically impossible) to capture this property of the data when fitting a statistical model. This separation step would also unduly blow up the number of observations resulting in a possible overestimation of our statistical certainty. For example, the 408 messages sent by actor 59 would result in 666 unicast messages after separation. Finally, note that directly modeling all possible compositions of the receiver set of multicast messages for a given sender in a network of N actors will become computationally infeasible as the number of compositions becomes tremedous (namely, $2^{N-1} - 1$ when excluding the empty receiver set). In sum, modeling email data (or relational data in general) with multicast messages is challenging but important for applied social network research.

To address this problem, this paper presents a multicast additive and multiplicative effects network (mc-amen) model for directly modeling unicast and multicast relational observations. The latent factor structure builds upon the amen model (Hoff (2015)). In the mc-amen model, for a given message of a given sender, all potential receivers are placed on a latent suitability scale where a higher (lower) suitability score makes it more (less) likely that the actor will be included in the receiver set in a message. Next, a threshold value is drawn from a truncated normal distribution that is bounded by the largest suitability score. Only the actors whose suitability score exceed the threshold value are included in the receiver set of the message. As will be shown, a truncated normal distribution for the threshold value results in the typical distributional shape of the number of receivers as in Figure 1. Moreover, the mc-amen model uses a multidimensional latent variable for the actors as well as for the messages to capture unobserved heterogeneity in the data. A multiplicative approach is considered due to its ability to capture unobserved interaction behavior between the actors, as implied by a latent distance model (Hoff, Raftery and Handcock (2002)) and as implied by a stochastic block model (DuBois, Butts and Smyth (2013), Nowicki and Snijders (2001)); for example, see Hoff (2008, 2009). As the interaction behavior of actors be be different as a sender and as a receiver, each actor has a the latent variable as a receiver and a latent variable as sender. When viewing each dimension of a latent dimension as a different topic (or working project),

these latent variables can be seen as a quantification of an actors' interest or involvement in this topic, either as a sender or as a receiver. Furthermore, the latent variables of the messages then quantify the degree of whether each topic is covered in a given message. Consequently, an actor is likely to become a receiver for a given message if the latent variables of this potential receiver, the latent variables of the sender, and the latent variables of the message are all large (in absolute manner) and have an equal sign. Additionally, the average popularity of actors is modeled using a random effect and the effects of known covariates are modeled using fixed effects. Finally, a normal error is added to the latent suitable scores, similar as a probit regression model, to allow a straightforward posterior computation using Gibbs sampling in a Bayesian framework.

The mc-amen model avoids potential limitations of related approaches for multicast messages. First, Perry and Wolfe (2013) extended the Cox proportional hazard model to relational observations using a partial likelihood formulation for the actors in a receiver set for a given sender, the observed time, and the number of receivers in a message. When conditioning on the cardinality of the receiver set, it is implicitly assumed that an actor first chooses the size of the receiver set and then chooses who to include as receivers in the message. This does not seem to be a realistic mechanism in real life applications. Furthermore, their approach does not provide a data generative mechanism for the number of receivers in a message, which is a key stochastic property in the data. Furthermore, their approach does not include any latent variables to capture potential unobserved heterogeneity. The mc-amen model on the other hand provides a data generative mechanism of both the size of the receiver set and the actors that are included in the receiver set. A flexible latent variable approach is implemented to capture unobserved heterogeneity regarding the interaction behavior between actors.

Second, Shafiei and Chipman (2010) proposed a mixed membership stochastic block model for the inclusion probabilities of the actors in the receiver set of a message sent by a given sender. In their latent variable approach, actors have the same latent variable as sender and receiver. This may not always be realistic, however. For example, a manager may be likely to receive updates from his/her team members about a certain topic or working project just to stay informed, but he/she may not be likely to send messages about this working project him/herself. Moreover, by not considering latent variables for the messages, the potential variability of topics in the messages is not captured. As will be shown in this paper, the latent variables of the messages are effective to capture the composition of the receiver sets in multicast messages. Another potential limitation of their approach is that the use of inclusion probabilities for each potential receiver results in a strictly positive probability of an empty receiver set. This contrasts with the observed data which only consists of receiver sets with at least one receiver. These limitations are automatically handled by the mc-amen model.

The paper is organized as follows. Section 2 presents the multiplicative latent factor model for relational data with multiple receivers. To fit the model, a Bayesian implementation is considered. Prior specification and posterior predictive checks for model assessment are discussed. Next, Section 3 describes numerical simulations to assess the behavior of the mc-amen model under different scenarios. Section 4 presents two analyses of the Enron email network. First, an empirical analysis is discussed of the subset of actors 26, 49, 53, 59, and 130 that was also discussed in this Introduction. Second, an empirical analysis of the full data set is discussed. The goal of the empirical analyses is to assess how well the mc-amen model can capture: (i) the empirical distribution of the cardinality of receiver sets of the entire data set (Figure 1) and (ii) the multicast interaction behavior between actors in the Enron network. We end the paper with a discussion in Section 5.

2. A multiplicative latent factor model for relational events with multiple receivers. We start with some notation. A network is considered for a set of actors denoted by $\mathcal{A}$ of size $|\mathcal{A}| = N$. The binary vector $\mathbf{y}_{ai}$ of length N captures the receiver set of the ith message

sent by actor a, where $y_{aia'} = 1$ if actor a' is in the receiver set, and $y_{aia'} = 0$ if a' is not in the receiver set, for $i = 1, \ldots, m_a$, where m_a is the number of messages sent by a. The total number of messages in the data is denoted by $M = \sum_{a=1}^{N} m_a$. The receiver set of message i of actor a is then written as $\mathcal{R}_{ai} = \{a' | y_{aia'} = 1\}$. We assume that an actor cannot send a message to one's self (even though the model would allow for this if needed). The reason is that messages sent to one's self do not contain any information about the relationships between actors.

We propose the following generative model for the binary receiver vector $\mathbf{y}_{ai}$ for message i sent by actor a:

$$\mathbf{z}_{ai,-a} \sim N(\boldsymbol{\theta}_{ai,-a}, \mathbf{I}_{N-1}),$$

(1)
$$c_{ai} | \mathbf{z}_{ai,-a} \sim TN(\mu_c, \sigma_c^2, -\infty, z_{ai(N)}),$$

$$y_{aia'} = \begin{cases} 1 & \text{if } z_{aia'} > c_{ai}, \\ 0 & \text{elsewhere,} \end{cases}$$

where the latent variable $z_{aia'}$ quantifies the suitability of message i for potential receiver actor a' sent by actor a on a latent suitability scale, $\theta_{aia'}$ is the mean suitability, c_{ai} is a threshold parameter for message i sent by a which follows a truncated normal distribution with mean μ_c and variance σ_c^2 in the subspace $(-\infty, z_{ai(N)})$, and where $z_{ai(N)}$ is the Nth ordered value of the latent suitability scores, that is, $z_{ai(N)} = \max_{a'}\{z_{aia'}\}$. By using a truncated distribution with upper bound $z_{ai(N)}$, the receiver set will never be empty. The mean μ_c and the variance σ_c^2 of the distribution of the threshold parameters control the distribution of the number of receivers across messages.

An additive and multiplicative latent effects model is considered for the mean suitability score according to

(2)
$$\theta_{aia'} = \mathbf{x}_{aia'}^{\top}\boldsymbol{\beta} + b_{a'} + \mathbf{u}_{a'}^{\top}\mathbf{v}_a + \mathbf{u}_{a'}^{\top}\mathbf{w}_{ai},$$

(3)
$$b_{a'} \sim N(\mu_b, \sigma_b^2),$$

where $\mathbf{x}_{aia'}$ is a vector of K observed predictor variables (e.g., exogenous or endogenous), $\boldsymbol{\beta}$ contains the K unknown coefficients that quantify the relative importance of these predictor variables, the random effect $b_{a'}$ quantifies the average popularity of actor a' as a receiver, which follows a normal distribution with mean μ_b and variance σ_b^2, $\mathbf{u}_{a'}$ is a vector containing the Q latent variables of actor a' as a receiver, $\mathbf{v}_a$ is a vector containing the Q latent variables of actor a as sender, and $\mathbf{w}_{ai}$ is a vector containing the Q latent variables of message i by sender a.

Each latent variable can be viewed as a theme or topic on which an actor can be (in)active as a sender or as a receiver or which is (not) important in a message. If a latent variable of actor a as sender, v_{aq}, is large in absolute value and of the same (opposite) sign as the latent variable of a potential receiver actor a', $u_{a'q}$, it becomes likely (unlikely) for the sender to include this potential receiver in the receiver set. Furthermore, the inner product of the latent vectors imply that the multiplicative function is added for the different latent dimensions, that is, $\mathbf{u}_{a'}^{\top}\mathbf{v}_a = u_{a'1}v_{a1} + \cdots + u_{a'Q}v_{aQ}$. Similarly, if a latent variable of message i, w_{aiq}, is large in absolute sense and of the same (opposite) sign as the latent variable of a potential receiver actor, $u_{a'q}$, it becomes likely (unlikely) to include this potential receiver in the receiver set of message i. By also including a latent variable for each message, the suitability scores of the potential receivers are not conditionally independent, given the latent variables of the sender and the potential receivers, the fixed effects, and the random popularity effects. Finally, note that by assuming different latent variables of each actor as a receiver and as a sender, the model allows for different roles of an actor as a receiver and as a sender on the latent topics.
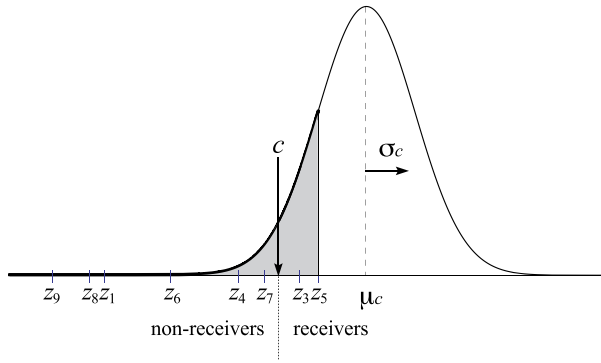
FIG. 3. *Graphical representation of the inclusion of the actors 3 and 5 in the receiver set for a message sent by actor 2 based on the random threshold value c. The index of the sender actor a and message i are omitted.*

Based on the fitted model to the data at hand, it can be investigated of course how strong the latent variables of the actors as senders and as receivers correlate with each other.

In sum, the model builds on the following intuition. When an actor decides to send a message, the sender actor determines the mean suitability for all potential receivers for the given message which is based on observed predictor variables (e.g., observed past interactions between actors or common locations where the actors work), the average popularity of the actors as a receiver, the multiplicative function of the latent variables of the sender and the potential receiver, and the multiplicative function of the latent variables of the message and the potential receiver. Random normal errors with a variance of 1 are added to the suitability scores. The suitability scores are placed on a latent scale. Subsequently, a threshold value is drawn from a truncated normal distribution where the upper bound equal to the largest suitability score. This ensures that the receiver set will always contain at least one actor, and the actor with the largest suitability score will always be included (Figure 3). All actors are included in the receiver set whose suitability score exceeds the drawn threshold value. Given the suitability scores of potential receivers, $\mathbf{z}_{ai,-a}$, the inclusion probability for actor a' in the receiver set $\mathcal{R}_{si}$ is then given by

$$\Pr(r \in R_{ai} | \mathbf{z}_{ai,-a}, \mu_c, \sigma_c) = \frac{\Phi(\frac{z_{aia'} - \mu_c}{\sigma_c})}{\Phi(\frac{z_{ai(N)} - \mu_c}{\sigma_c})},$$

where $\Phi(\cdot)$ denotes the standard normal cdf.

Because the mean of the threshold values, μ_c, and the mean of the popularity parameters, μ_b, both define the location of the latent suitability scale, we fix $\mu_c = 0$ and allow μ_b to be freely estimated to avoid identifiability issues. Moreover, the variance of the threshold distribution, σ_c^2, is difficult to identify when the means of the popularity parameters, μ_b, are freely estimated. This can be explained by the truncated normal distribution of the threshold parameter (having an increasing shape as in Figure 3), which remains virtually identical for specific combinations of the mean of the suitability scores and the standard deviation of the threshold distribution. To avoid any identifiability issues as a result of this, the standard deviation of the threshold distribution is fixed at $\sigma_c^2 = 1$.

Given this truncated normal distribution of the threshold parameter having a mean parameter of 0 and a variance of 1, the distribution of the cardinality of the receiver set will follow the typical decreasing shape we often see in email data with multicast messages (e.g., see Figure 1 for the Enron email data) when the means of the suitability scores are all negative (as in Figure 3). Consequently, the popularity parameters, b_a, which serve as intercepts for the mean of the suitability scores of each actor, will be negative when fitting the mc-amen model to such data.

2.1. *Prior specification and posterior computation.* For all free parameters, vague conjugate priors are considered to facilitate easy posterior sampling and to ensure that the posterior is completely dominated by the data,

$$\boldsymbol{\beta} \sim \mathcal{N}(\boldsymbol{\beta}_0, \boldsymbol{\Psi}_0),$$

$$\sigma_b^2 \sim IG(\alpha_b, \gamma_b),$$

$$\sigma_c^2 \sim IG(\alpha_c, \gamma_c),$$

$$\mathbf{v}_a, \mathbf{u}_{a'}, \mathbf{w}_{ai} \sim N(\mathbf{0}, \mathbf{D}),$$

where the hyperparameters are set to $\boldsymbol{\beta}_0 = \mathbf{0}$, $\boldsymbol{\Psi}_0 = 1\text{e}3\mathbf{I}_K$, $\alpha_b = \gamma_b = \alpha_c = \gamma_c = 1$,[1] and $\mathbf{D} = D\mathbf{I}_Q$. For the analysis of the Enron email data, a large prior variance is set for the latent variables of $D = 1\text{e}3$.

The conditional posterior distributions for $\boldsymbol{\beta}$, $\mathbf{b}$, $\mathbf{u}_r$, $\mathbf{v}_s$, and $\mathbf{w}_{si}$ then follow multivariate normal distributions, the random effects variance σ_b^2 follows inverse gamma distribution, and the threshold parameters c_{si} follow truncated normal distributions $N(\mu_c, \sigma_c^2)$ in the interval $(\max_{r:y_{sir}=0}\{z_{sir}\}, \min_{r:y_{sir}=1}\{z_{sir}\})$. Finally, the conditional posterior for the latent suitability scores is proportional to

$$\pi(\mathbf{z}_{si,-r} | \boldsymbol{\theta}_{si}, \mathbf{y}_{si,-r}, c_{si}) \propto N(\boldsymbol{\theta}_{si,-r}, \mathbf{I}_{P-1})$$

$$\times \prod_{r \neq s} 1(z_{sir} < c_{si})^{1-y_{sir}} 1(z_{sir} > c_{si})^{y_{sir}} \Phi\left(\frac{z_{si(P)} - \mu_c}{\sigma_c}\right)^{-1}.$$

To sample $\mathbf{z}_{si,-r}$, first a candidate $\mathbf{z}_{si,-r}^*$ is drawn from a multivariate truncated normal $N(\boldsymbol{\theta}_{si,-r}, \mathbf{I}_{P-1}) \prod_{r \neq s} 1(z_{sir} < c_{si})^{1-y_{sir}} 1(z_{sir} > c_{si})^{y_{sir}}$, and the draw is accepted with probability $\min(\frac{\Phi((z_{si(P)}^{(l-1)} - \mu_c)/\sigma_c)}{\Phi((z_{si(P)}^* - \mu_c)/\sigma_c)}, 1)$, where $\mathbf{z}_{si}^{(l-1)}$ is the previous ("$l - 1$"th) draw, using a Metropolis–Hastings step. These conditional distributions can be used to obtain a MCMC algorithm from which it is relatively straightforward to sample from, yielding posterior draws from the fitted model.

2.2. *Assessing model fit using posterior predictive checks.* Posterior predictive checks are flexible to assess whether a model captures certain key characteristics of the data (e.g., Gelman et al. (2004), Meng (1994), Van Kollenburg, Mulder and Vermunt (2015)). Moreover, in the presence of latent variables, posterior predictive checks are generally easier to apply (and perhaps easier to interpret) than traditional model fit indices, which assess global fit while correcting for model parsimony, such as the BIC. For this reason posterior predictive checks are used for evaluating the fit of the mc-amen model and, specifically, to determine the number of dimensions of the latent variable for the actors and messages. Replicated data sets are generated from the posterior predictive distribution using the posterior draws of the unknown parameters from the fitted mc-amen model. These replicated data sets have the same sample size as the observed data. Next, it is assessed whether certain characteristics of interest, referred to as test statistics, are similar between the replicated data and the observed data. If the differences are deemed to be too large, this suggests that the model badly captures this characteristic of the observed data.

The following test statistics are used for the posterior predictive checks of the mc-amen model:

---

[1]Based on our analyses, all results were virtually independent of the exact choice of the hyperparameters of the inverse gamma priors.

- The frequency of a specific combination of a sender and the receiver set in the observed data (either for unicast messages or multicast messages). This will be done for the most frequently observed combinations of senders and receiver sets.
- The distribution of the cardinality of the receiver sets over the entire data set. This will show whether the distributional form of the size of the receiver sets in the observed data is well-captured by the model, such as the typical decreasing trend as in the Enron data in Figure 1.
- The transitivity in the entire data set, given by

$$t_1(\mathbf{Y}) = \sum_{s \neq r} \sum_{j \in \mathcal{A} \setminus \{s,r\}} e_{s,j} e_{j,r} e_{s,r},$$

$$t_2(\mathbf{Y}) = \sum_{s \neq r} \sum_{j \in \mathcal{A} \setminus \{s,r\}} e_{s,j} e_{j,r} e_{r,s},$$

where $e_{s,r} = \sum_i y_{s,i,r} - \frac{1}{P-1} \sum_{i,r} y_{s,i,r}$, that is, the difference between the total number of messages sent by s to r receiver and the average number of receivers of messages sent by s. For the current paper, we consider a posterior predictive check for transitivity, which captures the tendency of actor s to send messages to actor r as a function of the number of messages sent by s to other actors than r and the number of messages sent by these other actors to r. Transitivity is often phrased as "the friends of my friends are my friends," an example of structural balance theory (Cartwright and Harary (1956), Heider (1946)). Transitivity can also be seen as a form of "broker-skipping," as it captures the tendency of actors to skip intermediate actors in a communication network (e.g., Leenders, Contractor and DeChurch (2016)). Multiplicative latent factor models, such as the mc-amen model, are known to capture such higher-order dependency structures in the data (Hoff (2005, 2009)).

Note that the first quantity, that is, the number of observed messages of a specific sender to a specific *receiver set*, is preferred over the number of messages of a specific sender to a specific receiver based on all receiver set for this given sender, that is, the relative popularity of an actor as a receiver for a given sender, when assessing model fit for relation with with multicast messages. For example, if actor 1 sent 100 unicast messages to actor 2, 100 unicast messages to actor 3, and no multicast messages to actors 2 and 3, the relative popularity of actors 2 and 3 would be the same for this sender as when actor 1 had send 100 multicast messages to both actors 2 and 3, and no unicast messages to actors 2 and 3 separately. Thus, by solely looking at the relative popularity, we would not be able to see whether the multicast behavior of actors in the observed data is well captured by the model.

3. Numerical simulations.

3.1. *Simulation* 1: *Parameter recovery and model fit under different populations.* Numerical simulations were performed to get insights about the recovery of the parameters and the performance of the model fit indicators. A network of 30 actors was considered having a one-dimensional latent variable. Table 1 (Population 0, first row) shows the parameter values in a reference population, and the parameter values for six other populations. In Populations 1a and 1b, the variances of the latent variable of the sender and the messages were different ($\text{var}(v_a) \ll \text{var}(w_{a,i})$ or $\text{var}(v_a) \gg \text{var}(w_{a,i})$). In Populations 2a and 2b, the correlations of the latent variables of actors as a sender and receiver differed ($\text{cor}(u_a, v_a) = 0.7$ or $\text{cor}(u_a, v_a) = -0.7$). In Population 3, the effects of the covariates was increased with a factor of 4 making the latent variables relatively less important. In Population 4 there were 10 out of 30 actors having a larger standard deviation for the threshold parameter (namely,

TABLE 1
Population values of the parameters for a network of 30 *actors with a one-dimensional latent variable. For empty cells the parameters are identical to the reference population (Population* 0, *first row). In Population* 4 *the standard deviation of the threshold parameter was* 1 *for* 20 *actors and* 5 *for* 10 *actors*

	β	(μ_b, σ_b)	$\text{var}(u_{1,a})$	$\text{var}(v_{1,a})$	$\text{cor}(u_{1,a}, v_{1,a})$	$\text{var}(w_{1,a,i})$	σ_c
Population 0	$(-\frac{1}{2}, -\frac{1}{4}, 0, \frac{1}{4}, \frac{1}{2})$	$(-7, 0.25)$	1	1	0	1	1
Population 1a				0.1		10	
Population 1b				10		0.1	
Population 2a					0.7		
Population 2b					-0.7		
Population 3	$(-2, -1, 0, 1, 2)$						
Population 4							1 or 5

$\sigma_c = 5$) resulting in considerably larger receiver sets for these 10 sending actors. Thus, under Population 4, the model is misspecified, as the variance of the threshold parameters is fixed at 1 for all actors under the mc-amen model. From each population, 25, 50, or 150 observations were generated from every sending actor. The covariates were generated from standard normal distributions. In the generated data, the proportion of multicast messages varied between 55% and 75% across the different populations. Models were fitted with $Q = 0$, 1, or a two-dimensional latent variable. To assess model fit, the posterior predictive distributions were considered for the frequencies of the 20 most commonly observed combinations of senders and receiver sets as well as the frequencies of the cardinality of the observed receiver sets. When fitting the models, we did not observe any label switching for the latent variables with for the model with a two-dimensional latent variable. Note that the labels of the two dimensions are arbitrary for this model type.

The 95% credibility intervals of the parameters and latent variables and posterior predictive distributions can be found in the Supplementary Material (Mulder and Hoff (2024)). The model with *zero*, one, and two latent variables are displayed in black, red, and green, respectively. For Population 0 to Population 4, Figures S.1–2 show the credibility intervals of the fixed effects, the random popularity parameters, and the latent variables of the actors as senders and receivers. Figures S.8–14 show the posterior means of the latent variables as a sender and as a receiver for the first latent dimension and for the second latent dimension (which is only available for the model with a two-dimensional latent variable). Finally, the posterior predictive distribution of the cardinality of the receiver sets and the frequencies of the 20 most commonly observed messages can be found in Figures S.15–21 and Figures S.22–28, respectively. Under Population 0 we can see that the parameters can be well-recovered (Figure S.1). Furthermore, we see that all three models (with $Q = 0$, one-, and two-dimensional latent variables) are able to capture the distribution of cardinality of the receiver set in the generated data (Figure S.15). Furthermore, Figure S.22 shows that the model with no latent variables cannot capture the observed frequencies of the 20 most commonly observed messages, while the model with a one-dimensional latent variable results in a satisfactory fit as well as the model with a two-dimensional latent variable shows hardly any improvement. Based on these plots, we would conclude that the model with a one-dimensional latent variable is preferred (as the fit is not improved for the mode with a two-dimensional latent variable).

Under Populations 1a and 1b, the results are very similar to the results under Population 0. Thus, different variances of the latent variables seem to be well-captured by the model (Figures S.19 and S.20). Under Populations 2a and 2b, the results are also comparable implying that positive and negative correlations between the latent variables can be well-captured by

the model. Under Population 3 all three models (with $Q = 0$, 1, or 2) fit the data approximately equally well (Figure S.23). This can be explained by the fact that the latent variables have a much smaller impact in comparison to the fixed effects. In this situation the model with zero latent variables would be preferred. Finally, under Population 4, the messages sent by the 10 actors with $\sigma_c = 5$ had receiver sets with an average size of 11 and a standard deviation of 10, and the messages sent by the other 20 actors with $\sigma_c = 1$ had receiver sets with an average size of 1.4 with a standard deviation of 0.8. Figure S.28 shows that the models with a zero- or one-dimensional latent variable result in a poor fit, which can be explained by model misspecification. Interestingly, however, the model with an additional latent dimension ($Q = 2$) results in an acceptable fit. This implies that the additional latent dimension is able to pick up on this specific characteristic of the data. This can also be seen from Figure S.14, which shows that the estimated latent variables of the actors as sender (v) are clustered for these 10 actors and for the other 20 actors for one dimension under the model with $Q = 2$. Furthermore, the model with a two-dimensional latent variable is also able to capture the distribution of the cardinality of the receiver sets, unlike the model with zero- or a one-dimensional latent variable (Figure S.21). In conclusion, these simulations show that the mc-amen model is able to capture a wide range of populations.

3.2. *Simulation* 2: *Simulations based on the Enron email network.* A simulation study was conducted to show the performance of the mc-amen model for analyzing network data with multicast messages and to determine the dimensionality of the latent variable using posterior predictive checks under a controlled situation. Network data were generated based on the characteristics of the Enron email data, which consisted of 156 actors who sent a total of 21,635 messages among each other varying from zero to 1597 with a median of 56 messages. To assess the effect of the sample size, three different sample sizes were considered in the simulation: the same number of messages as sent by the actors in the Enron data, half the number of messages as sent by the actors in the Enron data, and a quarter of the number of messages as sent by the actors in the Enron data. A two-dimensional latent variable was considered for the actors and messages where the latent variables of the actors as a receiver were generated using a bivariate normal distribution with zero means and a variances of 1.8 and a covariance of 0.9, the latent messages of the actors as sender were independently generated from a uniform distribution in the interval $(-1, 1)$, and the latent message variables were independently drawn from a discrete distribution having a value of 0 or 0.5 both with a probability of 0.5. These different distributions were considered to check how well the model can handle such different types of latent variables. Finally, the 32 random predictors were generated from normal distributions and the corresponding 32 coefficients (fixed effects) varied from -0.7 to 0.5 for β_1 to β_{32}, respectively, which is the same as the mc-amen model we will consider for the Enron email data (Section 4). Finally, the random popularity parameters were generated using a normal distribution with mean -8.7 and a standard deviation of 1. This resulted in a distribution of the cardinality of the receiver sets in the generated data that is similar as in the Enron data (Figure 1). Thereby, the simulation serves as a proof of concept for the analysis of the entire Enron email data.

For all three generated data sets, 600,000 posterior draws were obtained after burn-in of which every 100th draw was stored to correct for autocorrelations. For one generated sample, Figure 4 shows the posterior draws of the first five β coefficients and the five popularity parameters and latent variables as sender and receiver of the first five actors based on a fitted model with a three-dimensional latent variable for one data set (plots were made using the "bayesplot" package; Gabry and Mahr (2017)). Overall, the trace plots show convergence. The trace plots for the other parameters and other data sets looked similar.

Figure 5 shows the Bayesian credibility intervals of the 32 coefficients (top panels) and the 156 popularity parameters (lower panels, ordered by size) for a quarter of the sample size (left
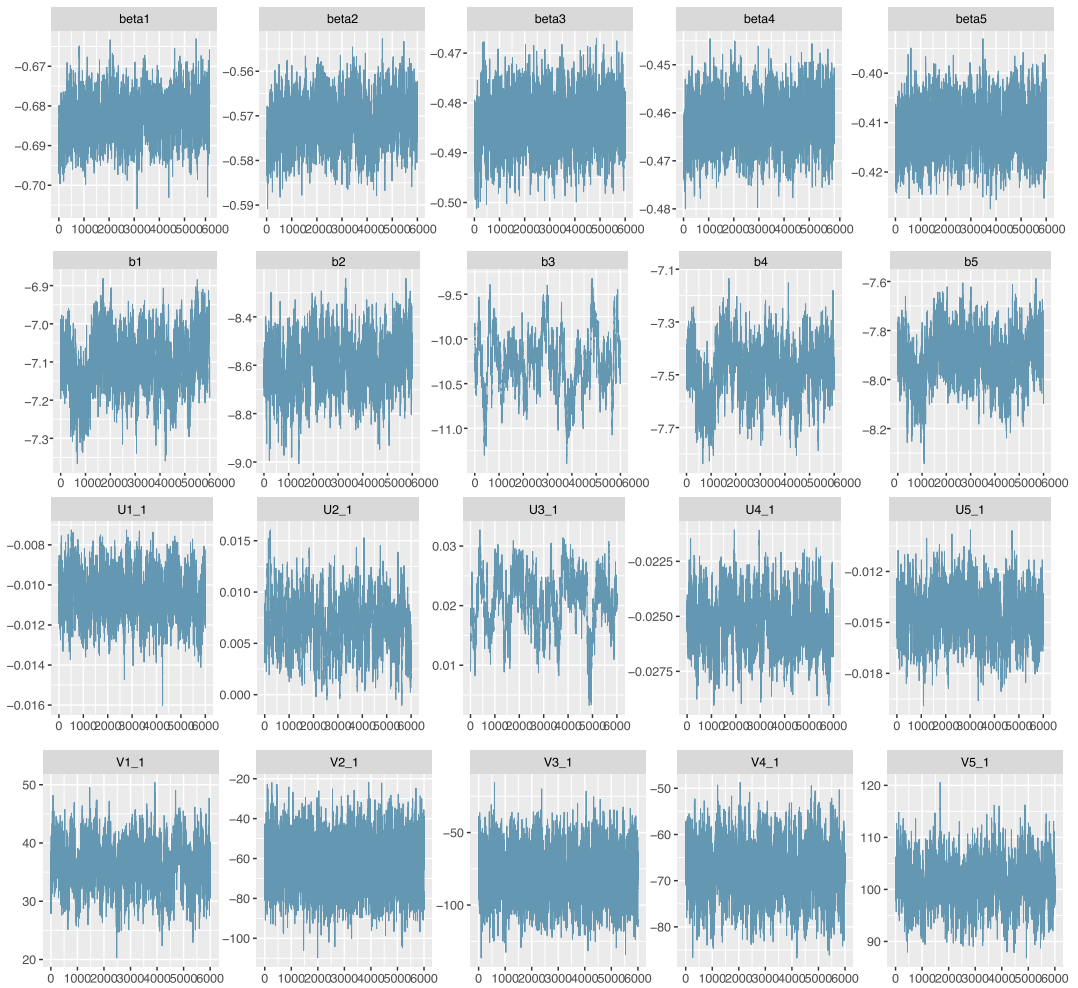
FIG. 4. *Trace plots of the posterior draws for a model with a latent variable with $Q = 0, 1, 2, and 3$ dimensions fitted to the simulated data.*

panels), half the sample size (middle panels), and the full sample size (right panels). For every parameter the intervals are displayed in batches of the fitted model with $Q = 0, 1, 2,$ and 3 dimensions for the latent variable, displayed in black, blue, purple, and red, respectively. Overall, the plots show that the relative magnitude of the parameters is fairly consistent for the different models with $Q = 0, 1, 2,$ and 3 dimensions of the latent variable. We see that for larger dimensions the popularity parameter becomes smaller on average. This can be a consequence of the variance of the suitability scores which increases for larger dimensions of the latent variable (because the mean suitability is modeled using model variance components). By accordingly decreasing the means of the suitability scores (which is controlled the popularity parameters), a similar distributional for the cardinality of the receiver sets can be obtained.

To assess model fit, predictive distributions of different statistical quantities were sampled under all fitted four models for all three data sets. To keep the presentation of the results as concise as possible, Figure 6 shows the posterior predictive distributions for the number of messages sent among the 40 most frequently observed combinations of senders and receivers for the three different sample sizes (quarter sample size in the panels in the first column, half the sample size in the second column, and the full sample size in the third column), the first transitivity for the full sample size (fourth column), and the distribution of the cardinality of
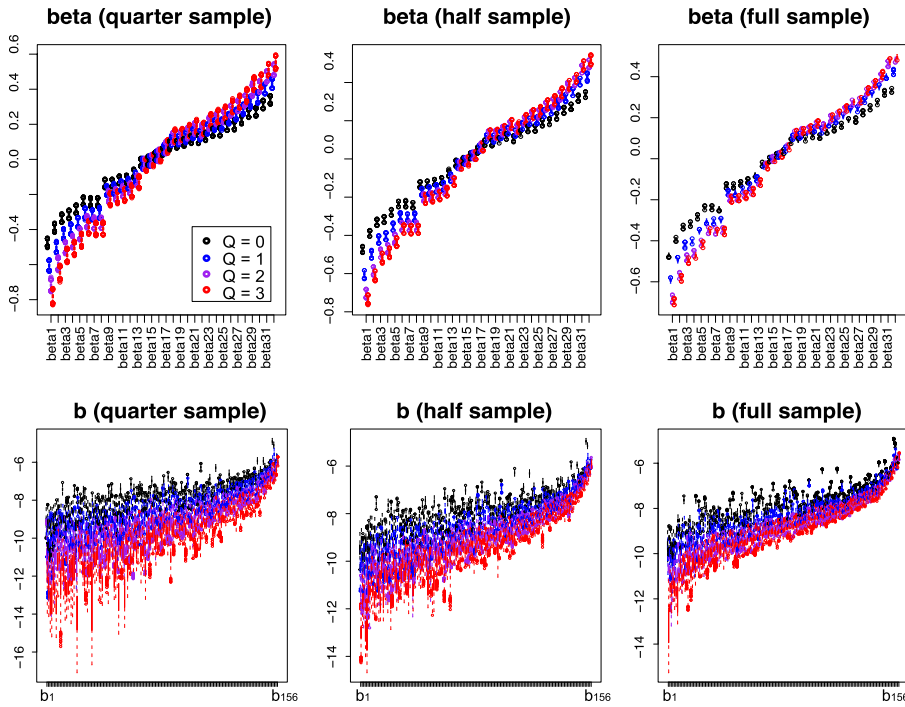
FIG. 5. *Posterior credibility intervals for the coefficients* $\beta_1, \ldots, \beta_{32}$ (*upper panels*) *and the popularity parameters* $b_1, \ldots, b_{156}$ (*lower panels*) *with a latent variable with* $Q = 0, 1, 2,$ *and* 3 *dimensions fitted to the simulated data for a quarter of the sample size* (*left panels*), *half the sample size* (*middle panels*), *and the full sample size* (*right panels*).

the receiver set in the full data (fifth column). The plots of the transitivity and the cardinality for half the sample size and the quarter of the sample size looked similar and are, therefore, omitted here. We see that the model without any latent variables (Figure 6, first row) does not capture the observed frequencies very well. The first transitivity statistic is also poorly captured by this model (fourth column). The distribution of the cardinality of the receiver sets however is well-captured by this model (fifth column), which illustrates that the mc-amen model is able to capture this property of the data very well, even when the model shows a poor fit based on the frequencies. We see an improvement of the fit of the mc-amen with a one-dimensional latent variable (second row), both in terms of the frequencies (in particular, the first few most frequently observed dyads) and in terms of transitivity. There is still some misfit though based on the observed frequencies for certain messages. This misfit seems no longer present when a mc-amen model with a two-dimensional latent variable is used (third row). This confirms our expectation because the data was also generated using a mc-amen with a two-dimensional latent variable. Model fit only improves very slightly more when further increasing the latent variable to $Q = 3$ dimensions (fourth row). Thus, to keep the model as parsimonious as possible, the mc-amen model with a two-dimensional latent variable seem to be preferred for all three data sets.

4. Analyses of the Enron email data.

4.1. *Analysis of a subset.* First we consider the subset of actors 26, 49, 53, 59, and 130 from the Enron email network discussed in the Introduction in this paper. The goal of this analysis is to see whether the mc-amen model is able to properly capture the multicast behav-
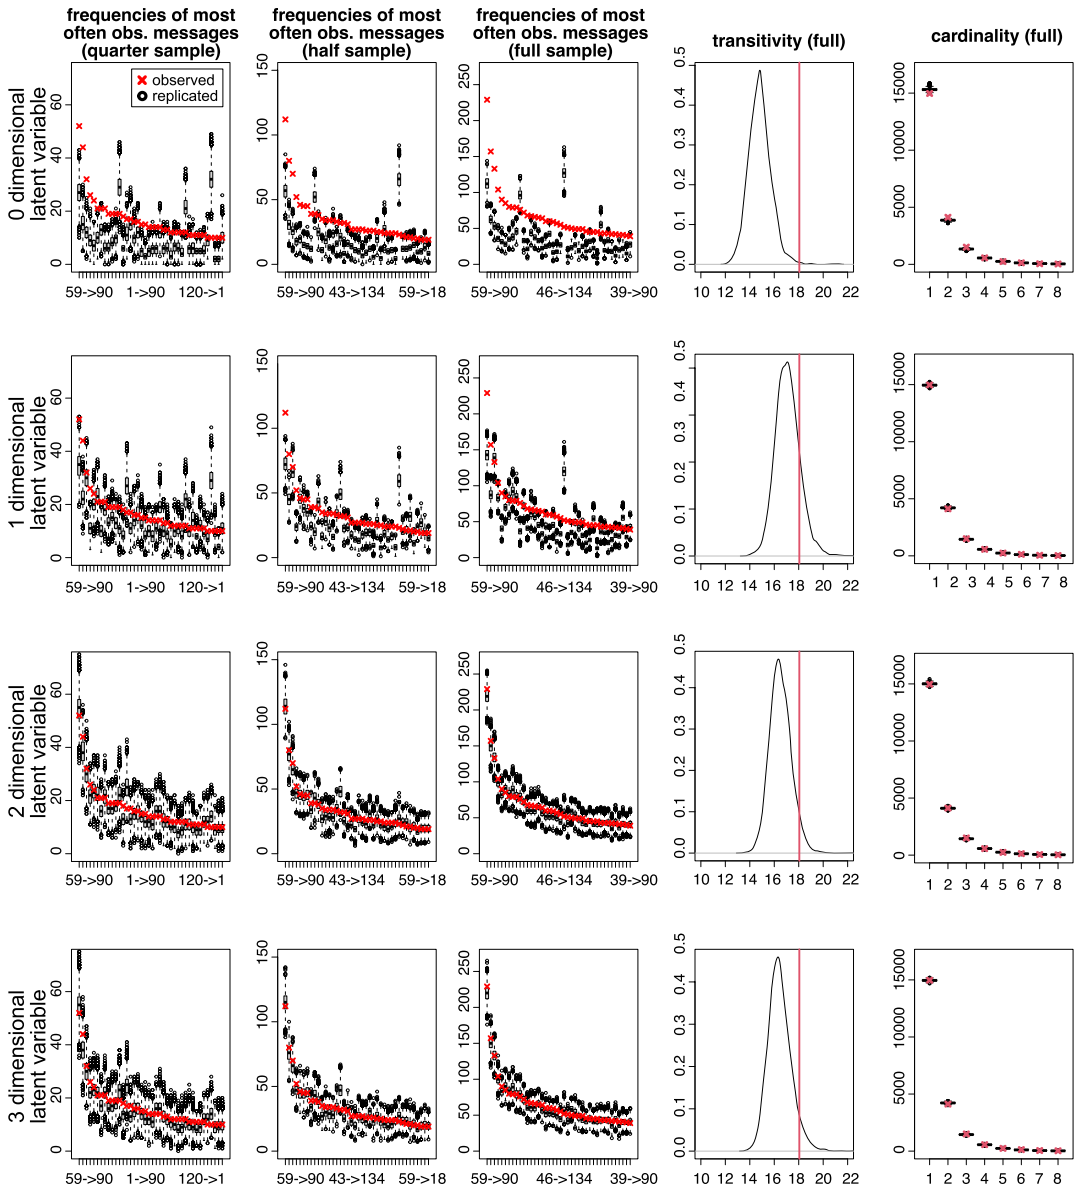
FIG. 6. *Posterior predictive distributions based on a model with a latent variable with $Q = 0, 1, 2,$ and 3 dimensions fitted to the simulated data. The red crosses and the vertical lines correspond to the respective quantities in the observed data.*

ior between the actors, to give some insights about the importance of latent variables to model multicast messages, and to show a comparison with existing approaches in a simple setup. No covariates were included in the analysis. In total, 624 were send among these actors with most messages (408) sent by actor 59 (see the receiver sets and in Figure 2). We fit a mc-amen model with no latent variables, a mc-amen model with a one-dimensional latent variable only for the actors and not for the messages (so without using $\mathbf{w}_{ai}$), a mc-amen model with a one-dimensional latent variable for all actors as well as messages, and a mc-amen model with a two-dimensional latent variable for all actors and messages. The first model with no latent variables is comparable with the model of Perry and Wolfe (2013), which also does not contain any latent variables. The second model with a one-dimensional latent variable for only the actors and not the messages is comparable with the model of Shafiei and Chipman (2010)
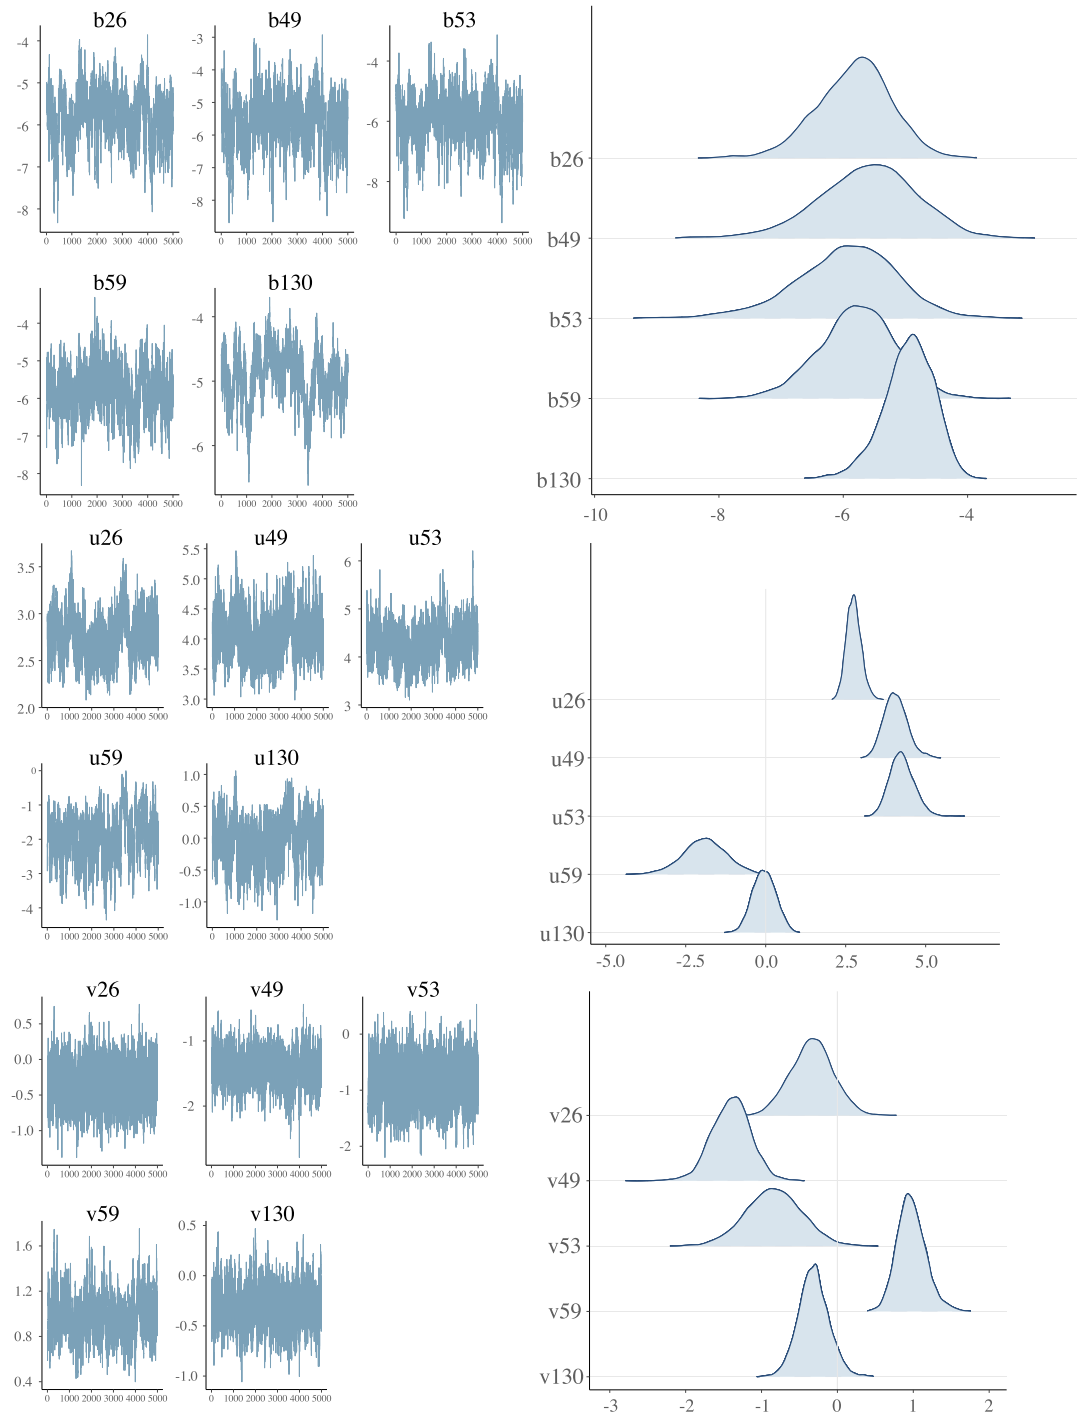
FIG. 7.   *Trace plots and posteriors distributions of the popularity parameters*, b_a, *and latent variables*, u_a *and* v_a, *of the five actors* $a = 26, 49, 53, 59$, *and* 130 *in the subset of the Enron network.*

(their model was more restrictive, however, because the latent variables of the actors were equal as sender and as receiver).

Figure 7 shows the trace plots of the posteriors of the popularity parameters and latent variables of the actors for the model with a one-dimensional latent variable for the actors and messages (the other models showed comparable results). Overall, the plots show reasonably
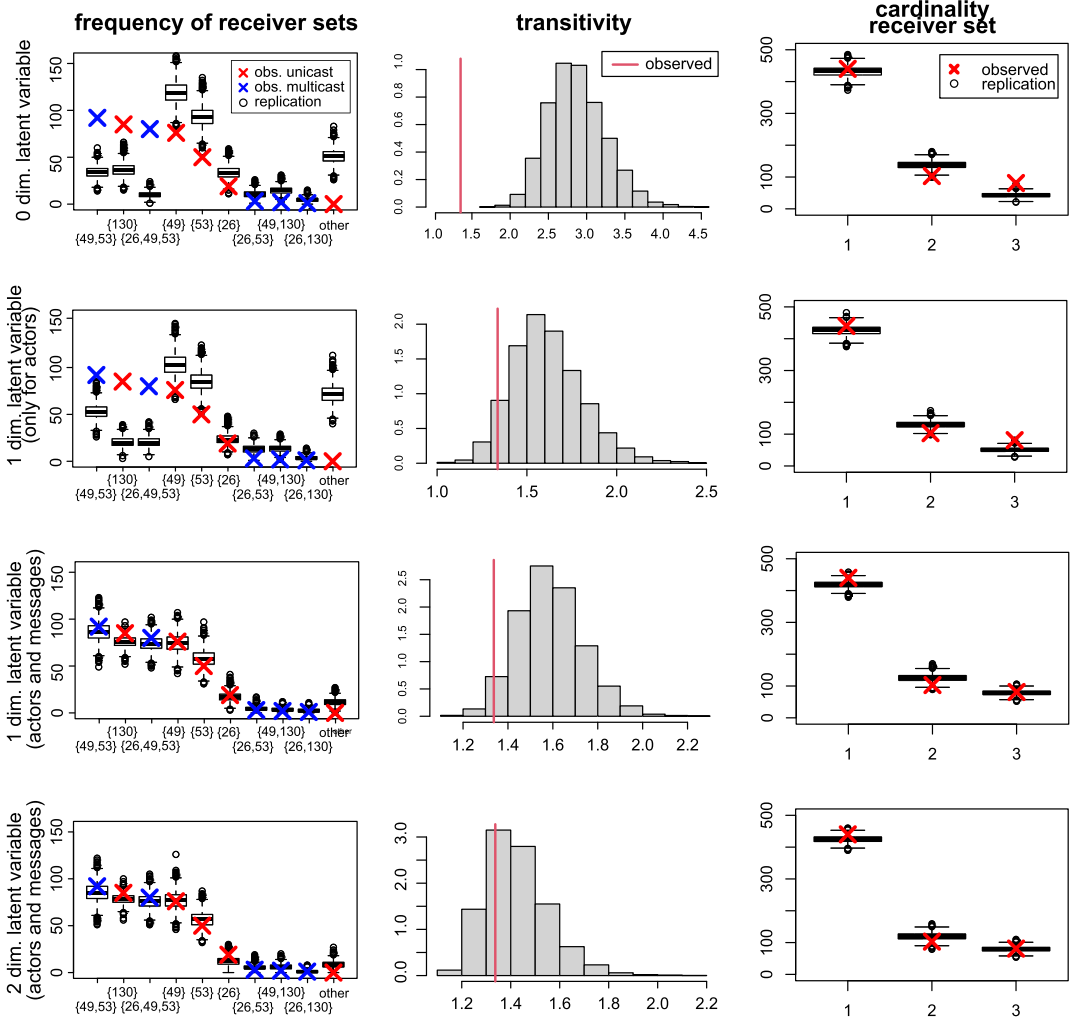
FIG. 8. *Posterior predictive distributions of the frequencies of the receiver sets of messages sent by actor* 59 (*left column*), *transitivity in the data* (*middle column*), *and the cardinality of the receiver set* (*right column*) *for a model with no latent variables* (*first row*), *a model with one latent variable with no latent variables for the messages* (*second row*), *a model with one latent variable for all actors and messages* (*third row*), *and a model with a two-dimensional latent variable for all actors and messages* (*fourth row*). *The crosses and vertical lines correspond to the respective quantities in the observed data.*

good mixing based on 500,000 draws of which every 100th draw is stored. Moreover, we see that the posteriors of the latent variables for the actors as receivers (u_a) differ from the latent variables as sender (v_a), which suggests that actors have different roles as sender and receiver regarding the unobserved heterogeneity.

Figure 8 (left column) shows the replicated frequencies for the nine observed receiver sets of messages sent by actor 59 and the remaining possible receivers sets (denoted by "other") using the four fitted models. We see that the first model with no latent variables and the second model with only a one-dimensional latent variable for the actors and not for the messages both badly fit the observed data, while the latter two models with latent variables for both the actors and messages fit the data quite well. Figure 8 (middle panels) shows the posterior predictive distributions of the first transitivity statistic for the four models (the plots for the second statistic looked comparable). We see that only the mc-amen model with no latent variables results in a poor fit while the other models fit this property of the data quite well, in
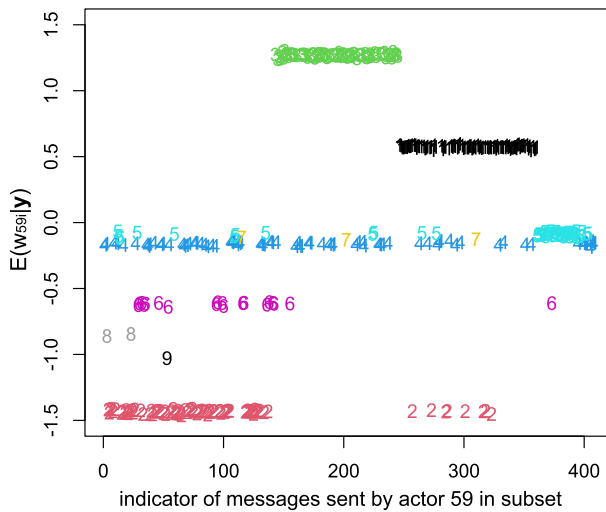
FIG. 9. *Posterior means of the latent variable of the* 408 *messages sent by actor* 59 *to actors* 26, 49, 53, *and* 130. *The numbers refer to the observed receiver sets of this sender from Figure* 8: *"1" =* {49, 53}, *"2" =* {130}, *"3" =* {26, 49, 53}, *"4" =* {49}, *"5" =* {53}, *"6" =* {26}, *"7" =* {26, 53}, *"8" =* {49, 130}, *and* *"9" =* {26, 130}.

particular, the model with a two-dimensional latent variable. Figure 8 (right column) shows that all four models capture the cardinality of the receiver sets over all messages in the data quite well. This again illustrates that the mc-amen model (also without any latent variables) is able to capture the observed distribution of the number of receivers per message in the complete data. Based on all these results, the third mc-amen model with a one-dimensional latent variable for the actors and messages seems to be preferred, as it results in an accurate fit to the data and the it is more parsimonious than the fourth model with a two-dimensional latent variable.

To give some insights of the importance of the latent variables of the messages to model multicast interaction behavior, Figure 9 displays the posterior means of the latent variables of the 408 messages sent by actor 59 where the nine different numbers (with different colors) refer to the nine different observed receiver sets for sender 59 (see left panels in Figure 8), for example, "1" refers to receiver set {49, 53}, "2" refers to receiver set {130}, etc. As can be seen the posterior means of the latent variables are practically equal for the same receiver sets. This shows that the latent variables of the messages play an important role when fitting the observed receiver sets. It is also interesting to see that for second receiver set "2," that is, {130}, the estimated latent variables are negative causing a large decrease of the suitability scores of actors 26, 49, and 53, due to their large and positive latent variables as receiver, that is, u_{26}, u_{49}, and u_{53} (Figure 7), and thus making only actor 130, with an estimated latent variable of approximately 0, to be a very suitable receiver. Similar explanations can be provided for the other receiver sets based on these results.

4.2. *Analysis of the full Enron email data.* The full Enron email data consisting of 21,635 messages which were sent among 156 actors between November 13, 1998, and June 21, 2002. The employees' genders (male, female), seniority (junior, senior), their type of work ("Legal Trading, Other"), title ('Specialist, Administrator," etc.) and their department ("ENA Gas Financial, Energy Operations," etc.) were also available. Following Perry & Wolfe (2013), we considered seven different inertia and reciprocity statistics for seven different time intervals, dummy covariates for all 16 combinations of senders and receivers working in "Legel" (1 = yes, 0 = no), working in "Trading" (1 = yes, 0 = no), being junior (1 = yes, 0 = no),
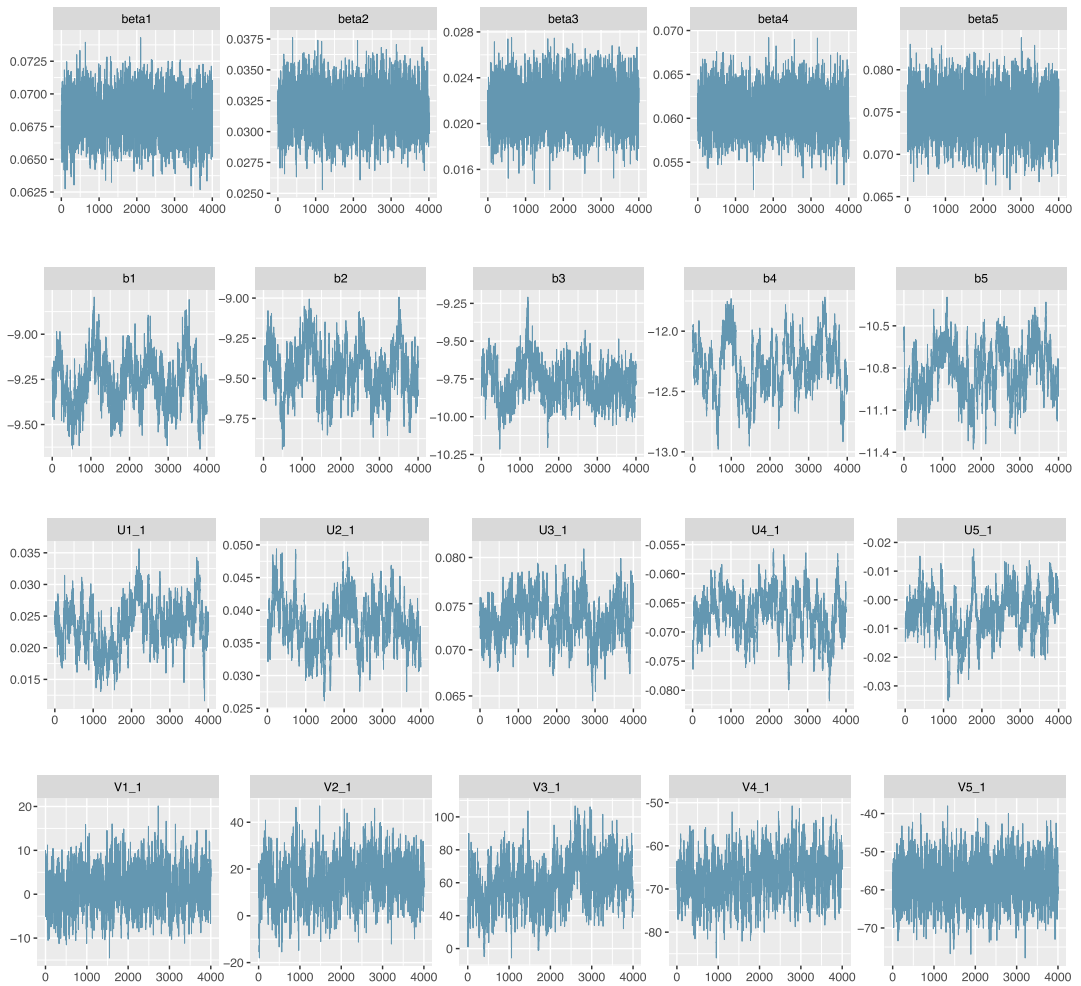
FIG. 10. *Trace plots of the first five coefficients* (*first row*), *the popularity parameters of the first five actors* (*second row*), *the first latent variable of the first five actors as receiver* (*third row*), *and the first latent variable of the first five actors as sender* (*fourth row*).

being female (1 = yes, 0 = no). We also added two dummy covariates of whether sender and receiver have the same title and whether they work in the same department. In total, this resulted in 32 covariates. All these covariates were standardized over the entire dataset.

A mc-amen was fit to these data with $Q = 0, 1, 2, 3$, and 4 dimensions. In the MCMC sampler, every 100th draw was stored. After burn-in, the trace plot of the posterior draws of the first five coefficients, the popularity parameters of the first five actors, the first latent variable of the first five actors as receiver, and the first latent variable of the first five actors as sender in the model with $Q = 2$ dimensional latent variable are displayed in Figure 10. The trace plots of the other parameters and other models were comparable and omitted to keep the results concise. Overall, the mixing looks acceptable.

Figure 11 shows the posterior credibility intervals of the coefficients (first row) and the popularity parameter (second row) for a mc-amen model with a latent variable with $Q = 0$, 1, 2, 3, and 4 dimensions fitted to the Enron email data. Again, we see that the relative size of the estimated coefficients is fairly consistent over all 32 parameters for the five different models with differing latent dimensionality. This suggests that conclusions about the relative importance of the covariates are largely unaffected of the dimensionality of the latent variable of the fitted mc-amen model. Furthermore, we see that the posterior intervals of the
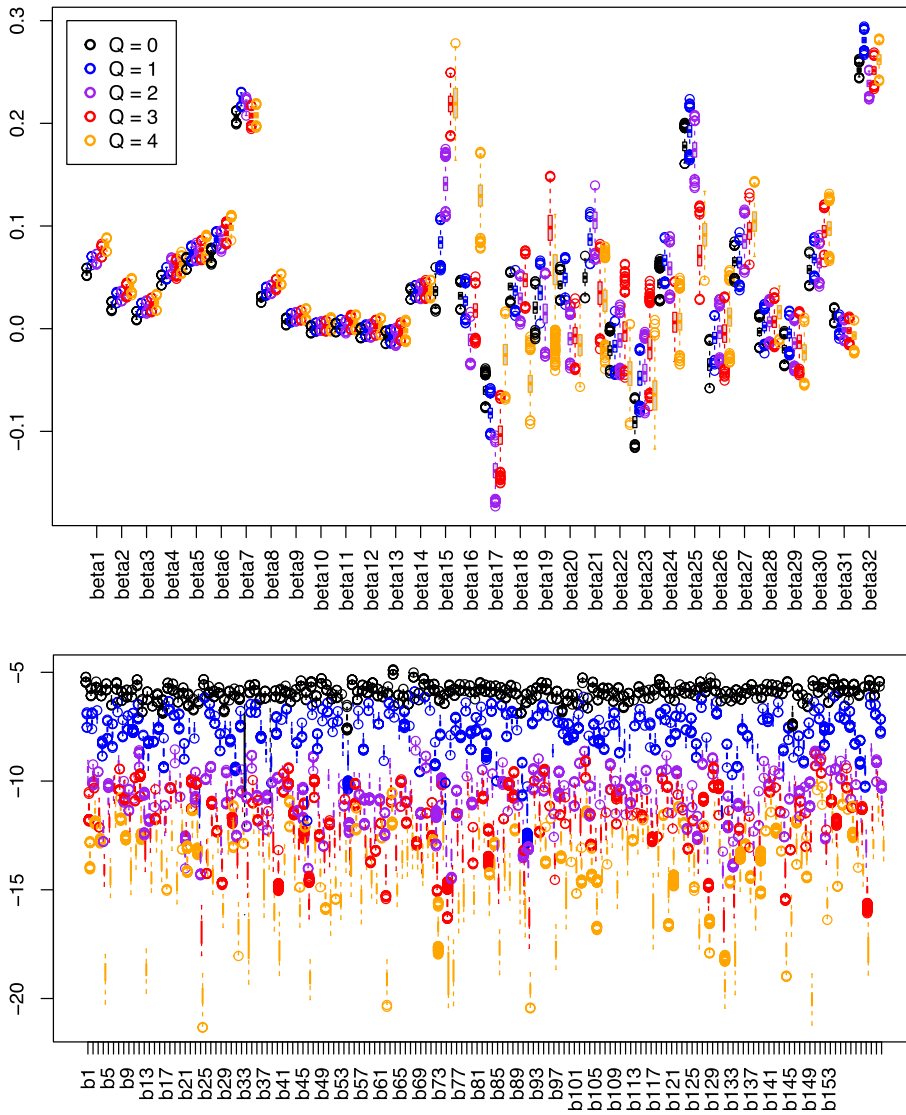
FIG. 11. *Posterior credibility intervals of the coefficients (upper panel) and the popularity parameter (lower panel) for a mc-amen model with a latent variable with Q = 0, 1, 2, 3, and 4 dimensions fitted to the Enron email data.*

popularity parameters become lower on average as the dimensionality increases, similar as in the simulated data.

Next, Figure 12 displays the posterior predictive distributions of the number of messages for the 20 most frequently observed combinations of senders and receivers in the observed data, the 21 to 80 most frequently observed combinations, the cardinality of the receiver sets over all data, and the two transitivity measures. In the first two columns the red and blue crosses denote the observed frequencies of the unicast and multicast messages, respectively. The model without latent variables (Figure 12, first row) clearly shows the worst fit. The model badly captures the frequency even of the most often observed messages (from actor 20 to actor 131) and also badly captures the frequencies of certain multicast messages (specifically of actor 20 to actors {61, 122, 131} and of actor 20 to actors {61, 113, 122, 131}). The misfit of the frequency of the number of messages of actor 20 to actor 131 can be explained by the fact that 131 is a popular receiver of messages sent by actor 20 for both unicast mes-
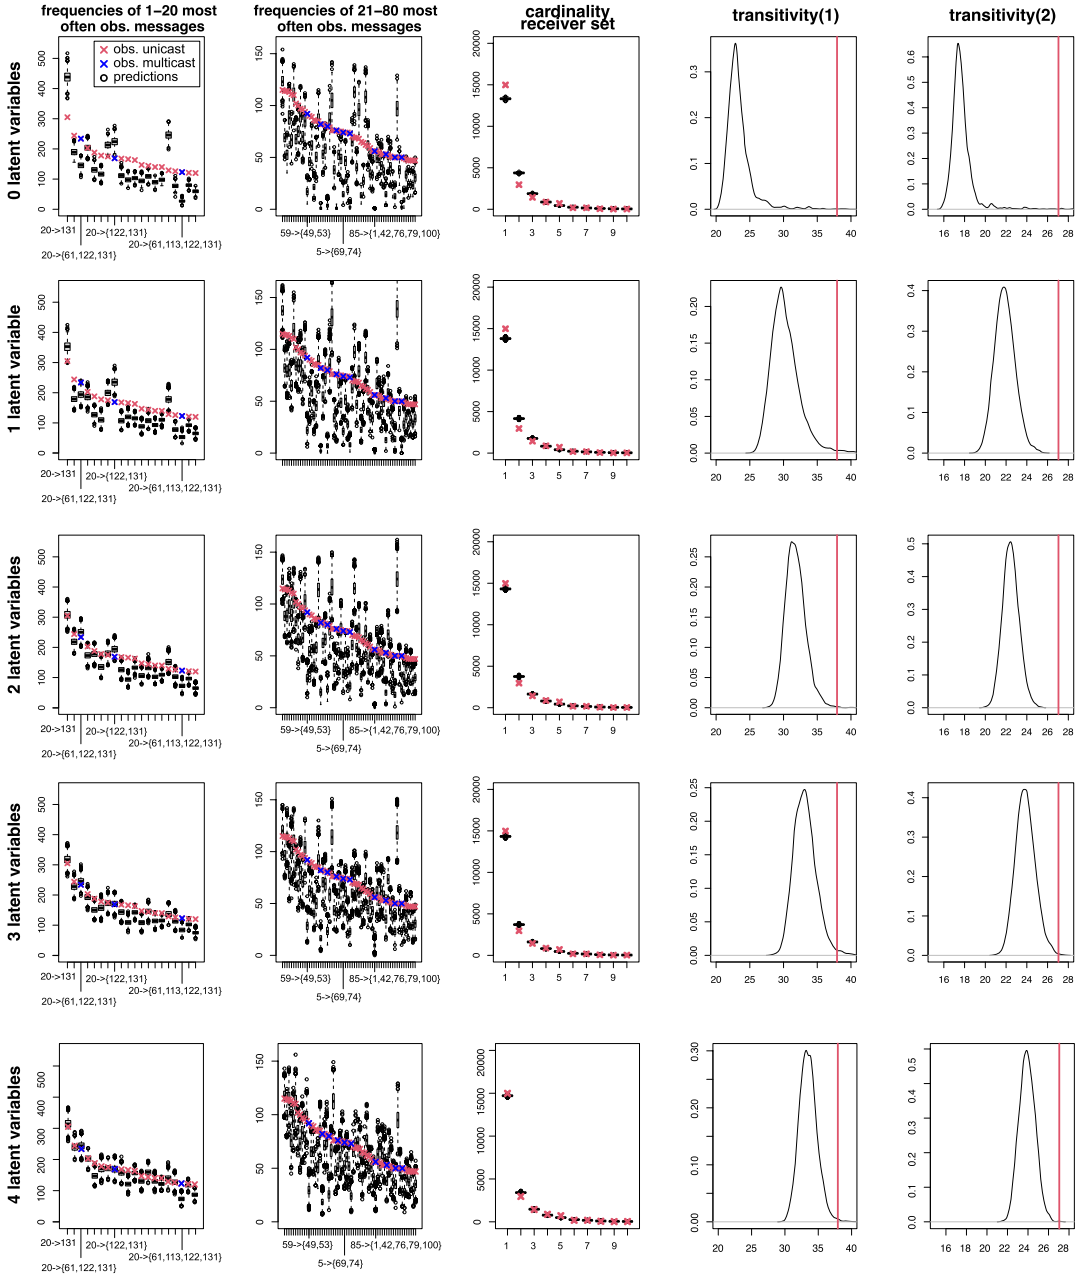
FIG. 12. *Posterior predictive distributions of the frequencies of the one to* 20 *most frequently observed messages* (*first column*), *the frequencies of the* 21 *to* 80 *most frequently observed messages* (*second column*), *the distribution of the cardinality of the receiver set over all messages* (*third column*), *the two transitivity measures* (*fourth and fifth column*) *for a model with zero, one, two, three, and four dimensions of the latent variable* (*first to fifth row*). *The crosses and vertical lines correspond to the respective quantities in the observed data.*

sages as well as multicast messages. Because the model assumes that unicast messages are on average more likely than multicast messages (because this is suggested from the data), the predicted frequencies of unicast messages of actor 20 to actor 131 become overestimated. This misfit considerably decreases for the mc-amen with a one-dimensional latent variable (Figure 12, second row), even though the frequencies of some of these 20 most commonly observed messages still show some misfit. This misfit has largely disappeared for the mc-amen
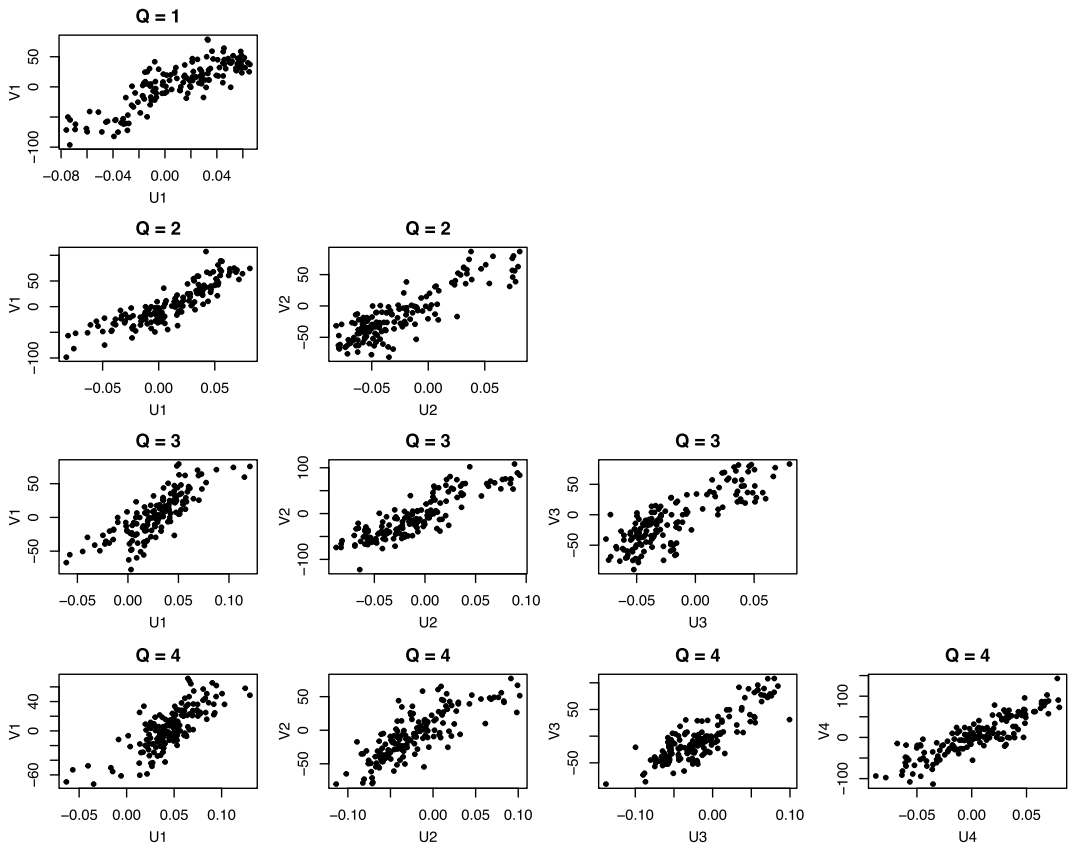
FIG. 13.   *Posterior means of each latent variable of an actor as receiver* (u_a) *vs. the latent variable as sender* (v_a) *for all latent variables in a mc-amen model with* $Q = 1, 2, 3,$ *and* 4 *dimensions for the latent variable.*

model with a two-dimensional latent variable (Figure 12, third row). For the 20 most frequently observed messages, the fit does not become much better for the mc-amen with a three-dimensional latent variable (Figure 12, fourth row) or a four-dimensional latent variable (Figure 12, fifth row). The same conclusion can be drawn when looking at the cardinality of the receiver set in the complete data set (Figure 12, third column) and the first transitivity measure (Figure 12, fourth column). For the second transitivity measure, the misfit may still be considered too large for the mc-amen with two dimensions. Based on this measure, a mc-amen model with three or four dimensions may be preferred. To avoid an overly complex model, while still maintaining a reasonable fit, the mc-amen model with a two-dimensional latent variable may be preferred based on these posterior predictive checks.

Next, we investigate how strong the latent variables of the actors as sender and as receiver are correlated. For this purpose the posterior means of latent variables under all models are computed. The results are displayed in Figure 13. The plots show the anticipated positive correlations for all latent variables under all models. Moreover, we see that the variability of the latent variables of actors as sender is considerably larger than the variability of the latent variable of actors as receiver. This, however, does not imply that actors vary more in their tendency to send messages in comparison to their tendency to receive messages because the fit is invariant to the scale of the latent variables, that is, $\mathbf{u}_{a'}^\top \mathbf{v}_a = (h\mathbf{u}_{a'})^\top (h^{-1}\mathbf{v}_a)$ and $\mathbf{u}_{a'}^\top \mathbf{w}_{ai} = (h\mathbf{u}_{a'})^\top (h^{-1}\mathbf{w}_{ai})$, for any constant $h > 0$. In the fitted model, the standard deviations of the estimated latent variables of the messages were approximately 30, which is comparable to the variability of the latent sender variables.

Based on these results, we can conclude the following. First, the mc-amen model is able to capture the (typical) decreasing shape of the empirical distribution of the cardinality of the receiver set in the Enron data (Figure 1) quite well. Second, the multiplicative latent factor structure in the mc-amen model is able to capture the multicast interaction behavior between actors quite well. Third, increasing the dimensionality of the latent variables generally result in an improved fit to the data. To avoid an overly complex model for the Enron data, the mc-amen model with a two-dimensional latent variable would be preferred. Fourth, the relative magnitude of the coefficients (fixed effects) depends little on the dimensionality of the latent variables.

5. Discussion. A multicast additive and multiplicative effects network (mc-amen) model was proposed for analyzing relational data that contain unicast messages (i.e., an observation of one actor toward another actor) and multicast messages (i.e., an observation of one actor toward multiple other actors). The multicast behavior is captured by placing the potential receivers of a given message sent by a given receiver on a suitability scale where a larger suitability score of a potential receiver results in a larger probability of being included in the receiver set. The suitability score is modeled using probit regression model of observed predictors variables and latent variables. Subsequently, a threshold parameter is drawn from a truncated normal distribution with an upper bound that is equal to the largest suitability score. Actors are included in the receiver set who have a larger suitability score than the threshold parameter. To capture any unobserved heterogeneity regarding the social interaction behavior in the network, a multiplicative approach is considered for the latent variables. This multiplicative approach encompasses both the discrete stochastic block model as well as the continuous latent distance model (Hoff (2008)). The mc-amen includes separate latent variables for the actors as sender and as potential receiver. Thereby, actors are allowed to have a different (latent) tendency to send messages as to receive messages. Furthermore, by including latent variables for the individual messages, it is possible to capture heterogeneous interaction behavior between a given sender and potential receivers.

The analyses of the Enron email data showed that the mc-amen model can accurately capture the empirical distribution of the cardinality of receiver sets that is typically observed in email data (where the probability tends to decrease as the cardinality increases, making unicast messages more likely than messages to two receivers, which in turn is more likely than messages to three receivers, etc.). Furthermore, the model is able to capture heterogeneous interaction behavior between actors in unicast messages as well as multicast messages. Furthermore, the data showed a positive correlation between the latent variables of the actors as receivers and as senders, which suggests that actors have a similar (but not identical) interaction behavior as a sender and as a receiver. To assess model fit posterior predictive checks were use based on the observed frequencies of the most commonly observed combinations of senders and receivers, the transitivity in the data, and the distribution of the cardinality of the receiver sets. This resulted in preference for a mc-amen model with a two-dimensional latent variable for the Enron data (where a more parsimonious model with fewer latent variables was preferred to avoid an overly complex model).

Different directions would be important to pursue for future research. In particular, it would be important to improve the computational efficiency of the algorithm when fitting the model. Because the mc-amen has many latent variables which depend on each other in a complex manner, posterior mixing can be slow resulting large autocorrelations (for this reason we used every 100th posterior draw in our analyses). Mixing could potentially be improved using orthogonal transformations of the parameters and by sampling the parameters in larger blocks where possible. Another potential computational improvement could be achieved using variational Bayesian algorithms. Another interesting direction for future research is to explore other distributions of the threshold parameters than the truncated normal

distributions. This is important for data with other types of shapes for the empirical distributions of the cardinality of the receiver set, for instance, when multicast messages are more frequently observed than unicast messages. A last interesting direction for future research we mention here is to extend the mc-amen model with an additional clustering step, which will allocate the actors in the same clusters if they show similar interaction behavior (e.g., similar as was done by Handcock, Raftery and Tantrum (2007), for the latent distance model). This could be done using a multivariate normal mixture model for the latent variables. Note that, even though the current mc-amen model does not parameterize this clustering step explicitly, it would of course be possible to identify clusters of actors if their estimated latent variables are close to each other based on the fitted mc-amen model.

SUPPLEMENTARY MATERIAL

Supplementary material (DOI: 10.1214/24-AOAS1885SUPP; .pdf). The Supplementary Material for this paper can be found in Mulder and Hoff (2024).

REFERENCES

ARENA, G., MULDER, J. and LEENDERS, R. T. A. (2022). A Bayesian semi-parametric approach for modeling memory decay in dynamic social networks. *Sociol. Methods Res.* 1–51.

BRANDES, U., LERNER, J. and SNIJDERS, T. A. B. (2009). Networks evolving step by step: Statistical analysis of dyadic event data. In 2009 *International Conference on Advances in Social Network Analysis and Mining* 200–205.

BUTTS, C. T. (2008). A relational event framework for social action. *Sociol. Method.* **38** 155–200.

CARTWRIGHT, D. and HARARY, F. (1956). Structural balance: A generalization of Heider's theory. *Psychol. Rev.* **63** 277–293. https://doi.org/10.1037/h0046049

COHEN, W. W. (2009). Enron email dataset.

DUBOIS, C., BUTTS, C. and SMYTH, P. (2013). Stochastic blockmodeling of relational event dynamics. In *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics* (C. M. Carvalho and P. Ravikumar, eds.). *Proceedings of Machine Learning Research* **31** 238–246. PMLR, Scottsdale, AZ, USA.

DUBOIS, C., BUTTS, C. T., MCFARLAND, D. and SMYTH, P. (2013). Hierarchical models for relational event sequences. *J. Math. Psych.* **57** 297–309. MR3137883 https://doi.org/10.1016/j.jmp.2013.04.001

ECKMANN, J.-P., MOSES, E. and SERGI, D. (2004). Entropy of dialogues creates coherent structures in e-mail traffic. *Proc. Natl. Acad. Sci. USA* **101** 14333–14337. MR2098979 https://doi.org/10.1073/pnas.0405728101

GABRY, J. and MAHR, T. (2017). bayesplot: Plotting for Bayesian models. R package version 1.

GELMAN, A., CARLIN, J. B., STERN, H. S. and RUBIN, D. B. (2004). *Bayesian Data Analysis*, 2nd ed. *Texts in Statistical Science Series*. CRC Press, Boca Raton, FL. MR2027492

HANDCOCK, M. S., RAFTERY, A. E. and TANTRUM, J. M. (2007). Model-based clustering for social networks. *J. Roy. Statist. Soc. Ser. A* **170** 301–354. MR2364300 https://doi.org/10.1111/j.1467-985X.2007.00471.x

HEIDER, F. (1946). Attitudes and cognitive organization. *J. Psychol.* **21** 107–112. https://doi.org/10.1080/00223980.1946.9917275

HOFF, P. D. (2005). Bilinear mixed-effects models for dyadic data. *J. Amer. Statist. Assoc.* **100** 286–295. MR2156838 https://doi.org/10.1198/016214504000001015

HOFF, P. D. (2008). Modeling homophily and stochastic equivalence in symmetric relational data. In *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics* (J. C. Platt, D. Koller, Y. Singer and S. Roweis, eds.). *Advances in Neural Information Processing Systems* **20** 657–664. MIT Press, Cambridge.

HOFF, P. D. (2009). Multiplicative latent factor models for description and prediction of social networks. *Comput. Math. Organ. Theory* **15** 261–272.

HOFF, P. D. (2015). Dyadic data analysis with amen. arXiv preprint. Available at arXiv:1506.08237.

HOFF, P. D., RAFTERY, A. E. and HANDCOCK, M. S. (2002). Latent space approaches to social network analysis. *J. Amer. Statist. Assoc*. **97** 1090–1098. MR1951262 https://doi.org/10.1198/016214502388618906

LEENDERS, R. T. A. J., CONTRACTOR, N. S. and DECHURCH, L. A. (2016). Once upon a time: Understanding team processes as relational event networks. *Organ. Psychol. Rev.* **6** 92–115. https://doi.org/10.1177/2041386615578312

MEIJERINK-BOSMAN, M., BACK, M., GEUKES, K., LEENDERS, R. and MULDER, J. (2022). Discovering trends of social interaction behavior over time: An introduction to relational event modeling: Trends of social interaction. *Behav. Res. Methods* 1–27.

MEIJERINK-BOSMAN, M., LEENDERS, R. and MULDER, J. (2022). Dynamic relational event modeling: Testing, exploring, and applying. *PLoS ONE* **17** e0272309. https://doi.org/10.1371/journal.pone.0272309

MENG, X.-L. (1994). Posterior predictive p-values. *Ann. Statist*. **22** 1142–1160. MR1311969 https://doi.org/10.1214/aos/1176325622

MULDER, J. and HOFF, P. D. (2024). Supplement to "A latent variable approach for modeling relational data with multiple receivers." https://doi.org/10.1214/24-AOAS1885SUPP

MULDER, J. and LEENDERS, R. T. A. J. (2019). Modeling the evolution of interaction behavior in social networks: A dynamic relational event approach for real-time analysis. *Chaos Solitons Fractals* **119** 73–85. MR3894308 https://doi.org/10.1016/j.chaos.2018.11.027

NOWICKI, K. and SNIJDERS, T. A. B. (2001). Estimation and prediction for stochastic blockstructures. *J. Amer. Statist. Assoc*. **96** 1077–1087. MR1947255 https://doi.org/10.1198/016214501753208735

PERRY, P. O. and WOLFE, P. J. (2013). Point process modelling for directed interaction networks. *J. R. Stat. Soc. Ser. B. Stat. Methodol*. **75** 821–849. MR3124793 https://doi.org/10.1111/rssb.12013

QUINTANE, E., CONALDI, G., TONELLATO, M. and LOMI, A. (2014). Modeling relational events: A case study on an open source software project. *Organ. Res. Methods* **17** 23–50. https://doi.org/10.1177/1094428113517007

SHAFIEI, M. and CHIPMAN, H. (2010). Mixed-membership stochastic block-models for transactional networks. In 2010 *IEEE International Conference on Data Mining*. https://doi.org/10.1109/ICDM.2010.88.

STADTFELD, C. and BLOCK, P. (2017). Interactions, actors, and time: Dynamic network actor models for relational events. *Sociol. Sci*. **4** 318–352. https://doi.org/10.15195/v4.a14

VAN KOLLENBURG, G. H., MULDER, J. and VERMUNT, J. K. (2015). Assessing model fit in latent class analysis when asymptotics do not hold. *Methodology* **1** 65–79.

ZHOU, Y., GOLDBERG, M., MAGDON-ISMAIL, M. and WALLACE, A. (2007). Strategies for cleaning organizational emails with an application to enron email dataset. In 5*th Conf. of North American Association for Computational Social and Organizational Science* **0621303**.