Marie Wiberg • Jee-Seon Kim • Heungsun Hwang •
Hao Wu • Tracy Sweet
Editors

Quantitative Psychology

The 88th Annual Meeting of the Psychometric Society, Maryland, USA, 2023

Springer

Editors
Marie Wiberg
Department of Statistics
Umeå School of Business,
Economics and Statistics
Umeå University
Umeå, Sweden

Heungsun Hwang
Department of Psychology
McGill University
Montreal, QC, Canada

Tracy Sweet
College of Education, HDQM
University of Maryland
College Park, MD, USA

Jee-Seon Kim
Department of Educational Psychology
University of Wisconsin-Madison
Madison, WI, USA

Hao Wu
Department of Psychology and Human
Development
Peabody College, Vanderbilt University
Nashville, TN, USA

ISSN 2194-1009 ISSN 2194-1017 (electronic)
Springer Proceedings in Mathematics & Statistics
ISBN 978-3-031-55547-3 ISBN 978-3-031-55548-0 (eBook)
https://doi.org/10.1007/978-3-031-55548-0

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Comparing Maximum Likelihood to Markov Chain Monte Carlo Estimation of the Multivariate Social Relations Model

Aditi M. Bhangale and Terrence D. Jorgensen

Abstract The social relations model (SRM) is a linear random-effects model applied to dyadic data within social networks (i.e., round-robin data). Such data have a unique nesting structure in that dyads (pairs) are cross-classified within individuals, who can also be nested in different networks. The SRM is used to examine basic multivariate relations between components of dyadic variables at two levels: individual-level random effects and dyad-level residuals. The current "gold standard" for estimating multivariate SRMs is the maximum likelihood (ML) estimation. However, Bayesian approaches, such as Markov chain Monte Carlo (MCMC) estimators, may provide some practical advantages to estimate complex or computationally intensive models. In this chapter, we report a small simulation study to compare the accuracy and efficiency of ML and MCMC point (and interval) estimates of a trivariate SRM on the ideal scenario: normally distributed, complete round-robin data. We found that MLE outperformed MCMC at both levels. MCMC greatly underestimated parameters and displayed poor coverage rates at the individual level but was relatively accurate at the dyad level.

1 Introduction

This chapter provides the first simulation study to compare the accuracy and efficiency of point and interval estimates of Markov chain Monte Carlo (MCMC) and maximum likelihood estimation (MLE) of multivariate social relations model (SRM) parameters. The SRM is a statistical and methodological approach traditionally applied to examine dyadic data gathered using a round-robin design (Gleason & Halperin, 1975). The round-robin design is typically a multiple-group reciprocal design wherein each group member interacts with or rates every other group member on some dyadic variable—for example, in group $g \in 1, \ldots, G$, member i reports

A. M. Bhangale (✉) · T. D. Jorgensen
University of Amsterdam, Department of Child Development and Education, Amsterdam, The Netherlands
e-mail: A.M.Bhangale@uva.nl; T.D.Jorgensen@uva.nl

M. Wiberg et al. (eds.), *Quantitative Psychology*, Springer Proceedings in Mathematics & Statistics 452, https://doi.org/10.1007/978-3-031-55548-0_7

their liking of member $j \neq i$—such that in a group of n_g individuals, each member participates in $n_g - 1$ pairs (dyads). Each interaction within a dyad $\{ij\}$ yields two observations—i's rating of j and j's rating of i—stored in a vector $\mathbf{y}_{\{ij\}}$. The braces {} indicate that when dyad members are indistinguishable (e.g., the same-sex peers are indistinguishable on the basis of sex), the order of i and j is arbitrary. Round-robin data have a unique nesting structure such that each dyadic observation $\mathbf{y}_{\{ij\}}$ is nested within both (a) data from dyads in which i is a member and (b) data from dyads in which j is a member. Thus, these designs allow decomposition of a dyadic variable into three[1] SRM components at two levels: out-going (ego) and in-coming (alter) effects at the individual level and dyadic (relationship) effects at the dyad level. In this manner, the SRM can quantify the degree to which the total variance in a dyadic variable is attributable to group- or individual-level differences versus the unique relationship shared between two individuals (Kenny et al., 2006, pp. 186–187).

1.1 The Univariate Social Relations Model

The following random-effects model (Gill & Swartz, 2001)

$$\mathbf{y}_{\{ij\}} = \begin{bmatrix} \mathbf{y}_{ij} \\ \mathbf{y}_{ji} \end{bmatrix} = \mu + \begin{bmatrix} E_i + A_j + R_{ij} \\ E_j + A_i + R_{ji} \end{bmatrix} \tag{1}$$

decomposes dyadic observations $\mathbf{y}_{\{ij\}}$ into individual-level ego (E) and alter (A) effects. E_i is an out-going effect representing, for example, how much i generally likes others. A_j is an in-coming effect indicative of how much j is generally liked by others. The relationship effect $R_{\{ij\}}$ is a residual effect composed of measurement error and i's unique liking of j beyond their individual tendencies to like others and be liked by others, respectively. Finally, μ is the grand mean of $\mathbf{y}_{\{ij\}}$ within the network (e.g., the average liking within a group).

Although individual-level effects are uncorrelated between individuals i and j, each individual's ego effect E_i and alter effect A_i are assumed to be bivariate normally distributed with expected value 0, variances σ_E^2 and σ_A^2, and a generalized covariance σ_{EA} (generalized reciprocity ρ_{EA} when standardized; Kenny et al., 2006, ch. 8):

$$\begin{bmatrix} E_i \\ A_i \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_E^2 & \\ \sigma_{EA} & \sigma_A^2 \end{bmatrix} \right). \tag{2}$$

[1] When sampling data from multiple round-robin groups, not only are dyads cross-classified within people, but people are also nested within groups, so group-level effects can also be decomposed from individual and relationship effects. However, estimation can be simplified by first partialing out group means (i.e., treated as fixed effects), simplifying the model by omitting group effects.

Positive σ_{EA} values indicate that if i generally likes others, then i is also generally liked by others. Negative σ_{EA} values indicate if i generally likes others, then i is generally *less liked* by others.

Likewise, R_{ij} and R_{ji} per dyad are assumed bivariate normally distributed:

$$\begin{bmatrix} R_{ij} \\ R_{ji} \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_R^2 & \\ \sigma_R^2 \rho_R & \sigma_R^2 \end{bmatrix} \right), \tag{3}$$

where relationship variances $\sigma_{R_{ij}}^2$ and $\sigma_{R_{ji}}^2$ are assumed to be equal ($\sigma_{R_{ij}}^2 = \sigma_{R_{ji}}^2 = \sigma_R^2$) when dyads are indistinguishable. The correlation between R_{ij} and R_{ji} effects per dyad is labeled the dyadic reciprocity ρ_R (Kenny et al., 2006, ch. 8). A positive ρ_R value means that if i particularly likes j, then j also particularly likes i beyond their individual-level tendencies to like others and be liked by others. A negative ρ_R value implies that if i particularly likes j, then j likes i particularly *less* than their individual-level tendencies to like others and be liked by others, respectively.

1.2 The Multivariate Social Relations Model

The SRM can be extended to multivariate cases. For example, Salazar Kämpf et al. (2018) investigated the association between liking of strangers at first impression and subsequent mimicry during a 5-minute interaction. In this scenario, the vector of SRM equations expands as follows:

$$\begin{bmatrix} \mathbf{y}_{\{ij\}} \\ \mathbf{z}_{\{ij\}} \end{bmatrix} = \begin{bmatrix} \mathbf{y}_{ij} \\ \mathbf{y}_{ji} \\ \mathbf{z}_{ij} \\ \mathbf{z}_{ji} \end{bmatrix} = \begin{bmatrix} E_{\mathbf{y},i} \\ E_{\mathbf{y},j} \\ E_{\mathbf{z},i} \\ E_{\mathbf{z},j} \end{bmatrix} + \begin{bmatrix} A_{\mathbf{y},j} \\ A_{\mathbf{y},i} \\ A_{\mathbf{z},j} \\ A_{\mathbf{z},i} \end{bmatrix} + \begin{bmatrix} R_{\mathbf{y},ij} \\ R_{\mathbf{y},ji} \\ R_{\mathbf{z},ij} \\ R_{\mathbf{z},ji} \end{bmatrix}, \tag{4}$$

where $\mathbf{y}$ are pre-interaction liking ratings and $\mathbf{z}$ are mimicry ratings.

Similar to the univariate case, individual-level effects of every person are assumed to be multivariate normally distributed so that individual-level covariances between the two dyadic variables can be estimated:

$$\begin{bmatrix} E_{\mathbf{y}_i} \\ A_{\mathbf{y}_i} \\ E_{\mathbf{z}_i} \\ A_{\mathbf{z}_i} \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_{E_\mathbf{y}}^2 & & & \\ \sigma_{A_\mathbf{y},E_\mathbf{y}} & \sigma_{A_\mathbf{y}}^2 & & \\ \sigma_{E_\mathbf{z},E_\mathbf{y}} & \sigma_{E_\mathbf{z},A_\mathbf{y}} & \sigma_{E_\mathbf{z}}^2 & \\ \sigma_{A_\mathbf{z},E_\mathbf{y}} & \sigma_{A_\mathbf{z},A_\mathbf{y}} & \sigma_{A_\mathbf{z},E_\mathbf{z}} & \sigma_{A_\mathbf{z}}^2 \end{bmatrix} \right). \tag{5}$$

For example, $\sigma_{A_\mathbf{z},E_\mathbf{y}}$ is an ego–alter covariance that can be used to investigate whether individuals who are generally liked more at first impression also display greater mimicry during a 5-minute interaction. Likewise, ego–ego covariances (e.g., $\sigma_{E_\mathbf{y},E_\mathbf{z}}$) and alter–alter covariances (e.g., $\sigma_{A_\mathbf{y},A_\mathbf{z}}$) can be estimated.

At the dyad level, the interchangeability of $\{ij\}$ implies constraints:

$$\begin{bmatrix} R_{\mathbf{y}_{ij}} \\ R_{\mathbf{y}_{ji}} \\ R_{\mathbf{z}_{ij}} \\ R_{\mathbf{z}_{ji}} \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma^2_{R_\mathbf{y}} & & & \\ \sigma^2_{R_\mathbf{y}} \rho_{R_\mathbf{y}} & \sigma^2_{R_\mathbf{y}} & & \\ \sigma^{\text{intra}}_{R_\mathbf{z}, R_\mathbf{y}} & \sigma^{\text{inter}}_{R_\mathbf{z}, R_\mathbf{y}} & \sigma^2_{R_\mathbf{z}} & \\ \sigma^{\text{inter}}_{R_\mathbf{z}, R_\mathbf{y}} & \sigma^{\text{intra}}_{R_\mathbf{z}, R_\mathbf{y}} & \sigma^2_{R_\mathbf{z}} \rho_{R_\mathbf{z}} & \sigma^2_{R_\mathbf{z}} \end{bmatrix} \right). \tag{6}$$

Note the equality-constrained intrapersonal covariance (e.g., $\sigma^{\text{intra}}_{R_\mathbf{z}, R_\mathbf{y}}$), which estimates the degree to which i's unique liking of j at the first impression covaries with i's unique mimicry of j. Note also the equality-constrained interpersonal covariance (e.g., $\sigma^{\text{inter}}_{R_\mathbf{z}, R_\mathbf{y}}$), which estimates the degree to which i's unique liking of j at the first impression covaries with j's unique mimicry of i. These equalities follow from dyad members being indistinguishable, similar to the equality-constrained relationship variance per variable.

1.3 Estimation of Social Relations Models

Several techniques including ANOVA-based estimators (Warner et al., 1979), MLE, and MCMC have previously been applied to estimate the SRM parameters introduced in the previous sections. For the purpose of this chapter, we concentrate on MLE and MCMC.

The MLE approach (Nestler, 2018; Nestler et al., 2020), applied in the R package `srm` (Nestler et al., 2022), is the most recently proposed method for estimating SRM parameters. Nestler et al. (2020) apply a Fisher-scoring algorithm to derive SRM (co)variance point and SE estimates, which is applicable with unbalanced or incomplete normally-distributed data. However, the accuracy of this method generally depends on the sample size (Hoff, 2005) and the shape of the SRM variance components' sampling distributions (Lüdtke et al., 2013).

MCMC estimators, for example, Gibbs sampling (Gill & Swartz, 2001; Hoff, 2005; Lüdtke et al., 2013), provide some practical advantages. The specification of prior distributions can incorporate expectations and previous knowledge about SRM parameters. Point estimators of SRM (co)variances may be selected as the mean (expected *a posteriori*, EAP), median, or mode (maximum *a posteriori*, MAP) of the empirical posterior distribution. These approaches handle unbalanced and incomplete designs (Gill & Swartz, 2001) and can incorporate model uncertainty well. Additionally, Bayesian estimators possess the added benefit of estimating complex or computationally intensive models that MLE cannot. However, the posterior summaries may perform differently depending on the nature of the data (see Lüdtke et al., 2013).

In this chapter, we compared MLE and MCMC estimation of (co)variances of trivariate SRM components.

2 Method

2.1 Population Values and Simulation Conditions

We used population (co)variance matrices from the Open Science Framework (OSF) project of Nestler et al. (2020, see https://osf.io/9twkm/):

$$\Sigma_{EA} = \begin{pmatrix} 0.600 \\ 0.480 & 0.776 \\ 0.280 & 0.336 & 0.396 \\ 0.100 & 0.060 & 0.035 & 0.300 \\ 0.030 & 0.036 & 0.021 & 0.120 & 0.172 \\ 0.030 & 0.036 & 0.051 & 0.120 & 0.072 & 0.172 \end{pmatrix} \text{ and}$$

$$\Sigma_{R} = \begin{pmatrix} 0.900 \\ 0.150 & 0.900 \\ 0.480 & 0.120 & 0.884 \\ 0.120 & 0.480 & 0.196 & 0.884 \\ 0.840 & 0.210 & 0.672 & 0.168 & 1.576 \\ 0.210 & 0.840 & 0.168 & 0.672 & 0.094 & 1.576 \end{pmatrix}.$$

The rows and columns of Σ_{EA} are arranged such that the first three rows correspond to the ego effects and the last three rows correspond to the alter effects of the dyadic variables. The rows and columns of Σ_R are arranged such that the first two rows correspond to the ij and ji effects of the first dyadic variable, the center two rows correspond to the ij and ji effects of the second dyadic variable, and the final two rows correspond to the ij and ji effects of the third dyadic variable.

Round-robin group size (n_g) and the number of round-robin groups (G) were manipulated such that as n_g and G increased, so did the number of persons and dyads within the sample. We sampled $G = 10$ or 20 networks of size $n_g = 6$, 8, or 10. We generated 100 samples per sample size condition.

2.2 Analysis Plan

All analyses were conducted in R (R Core Team, 2023). Specifically, the SRM (co)variances for each of the 100 samples across the six conditions were estimated with MLE using the `srm` package (Nestler et al., 2022) and with a modified Hamiltonian Monte Carlo (HMC) algorithm called the No-U-Turn Sampler (NUTS; Hoffman et al., 2014), which is available in the `rstan` package (Stan Development Team, 2023).

The `srm` package assumes the SRM components to be multivariate normally distributed latent factors with mean vector μ and covariance matrix Σ and can accommodate structural relations between SRM components. However, we fit a saturated model at each level, which results in unconstrained (co)variance estimates between the SRM components (other than equality constraints for indistinguishable dyads), which is simply a multivariate SRM.

The `mvsrm()` function within the `lavaan.srm` package estimates a multivariate SRM using MCMC estimation via `rstan`. The `lavaan.srm` package specifies diffuse prior distributions for all parameters by default. Priors for the *SD*s of level-specific effects were student-*t* distributions:

$$\sigma_{EA(or R)} \sim t(\nu = 4, \mu = 0.5, \sigma = 0.5). \tag{7}$$

The correlation matrix at the individual level followed an LKJ distribution (Lewandowski et al., 2009):

$$R_{EA} \sim LKJ(\eta = 2), \tag{8}$$

whereas each nonredundant dyad-level correlation followed a beta distribution:

$$R_R \sim Beta(\alpha = 1.5, \beta = 1.5), \tag{9}$$

from which sampled parameters were rescaled as $2x - 1$ to provide support across the range $\{-1, +1\}$ rather than the usual $\{0, 1\}$ range.

Each person's vector of random effects (E and A per variable) was sampled from a multivariate standard normal distribution (sampling z scores is more computationally stable). Sampled random effects $\widehat{E}$ and $\widehat{A}$ were scaled by their estimated *SD*s when calculating expected values per dyad and per variable:

$$\widehat{y}_{ij} = \widehat{\sigma}_E \widehat{E}_i + \widehat{\sigma}_A \widehat{A}_j. \tag{10}$$

Finally, the likelihood for the round-robin observations y_{ij} followed a multivariate normal distribution with mean equal to each dyad's expected values and covariance matrix equal to $\widehat{\Sigma}_R$:

$$y_{ij} \sim \mathcal{MVN}(\mu = \widehat{y}_{ij}, \Sigma = \widehat{\Sigma}_R). \tag{11}$$

We initialized four Markov chains with random starting values, running each for 2000 iterations and discarding the first half of each as burn-in samples. This yielded 4000 posterior samples to estimate the joint posterior distribution of SRM parameters. We monitored the bulk effective sample size (ESS) and $\widehat{R}$ to check for convergence issues. If ESS < 100 or $\widehat{R} > 1.05$, we repeated MCMC estimation with double the iterations (again discarding the first half as burn-in). For each sampled data set, we saved the EAP and MAP estimates of SRM (co)variances, the latter being analogous to ML estimates.

In both `srm` and `lavaan.srm`, group effects were treated as fixed by group-mean centering all dyadic variables. We estimated 33 unique (co)variances across the two levels: 21 at the individual level and 12 at the dyad level. To evaluate the accuracy of each estimator consistently with Nestler et al. (2020), we inspected the robust relative bias (RB) of the point estimates and the coverage rates (CRs) of the interval estimates. We also computed the relative efficiency (RE) with respect to the root mean-squared error (RMSE) of the estimators.

3 Results

The original four Markov chains appeared to converge on the same posterior distribution within 2000 iterations, for 98–100% of all samples per condition. For seven samples that underwent 4000 iterations, ESS > 100 or $\widehat{R} < 1.05$ indicated that doubling the number of posterior samples appeared to resolve convergence issues. The results for all 33 (co)variances are visualized in plots presented in Figs. 1, 2, 3, 4, 5 and 6 and are summarized below. All the results can be found in our supplementary material on the OSF: https://osf.io/w3jue/.

3.1 *Accuracy of Point Estimates*

As shown in Fig. 1, parameters are, on average, underestimated at the individual level. Some ML estimates display acceptable bias; however, most parameters are greatly underestimated across all conditions. The bias in EAP and MAP estimates is even greater. When estimated via MCMC, the RBs of two parameters are greatly exacerbated due to dividing by near-zero values: the generalized covariance of the second dyadic variable ($\sigma^2_{E_1 A_1}$, population value = 0.036) and the covariance between the alter effect of the second dyadic variable and ego effect of the third dyadic variable ($\sigma^2_{A_2 E_3}$, population value = 0.021).

In contrast with the individual level, most estimates at the dyad level are overestimated (see Fig. 2). However, relative biases have much smaller magnitude than at the individual level, indicating that parameters are more accurately estimated at the dyad level. One reason for this is that whereas information from only $n_g \times G$ individuals is used to estimate individual-level parameters, information from

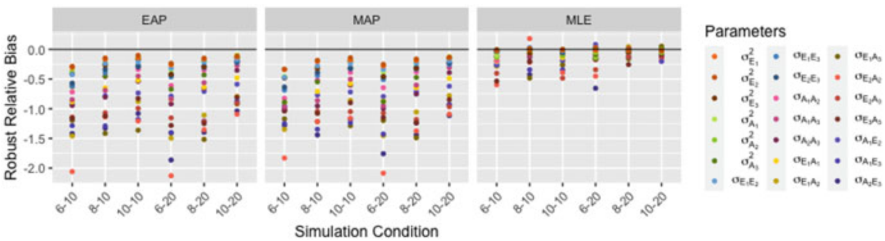

Fig. 1 Robust relative bias of point estimates at the individual level

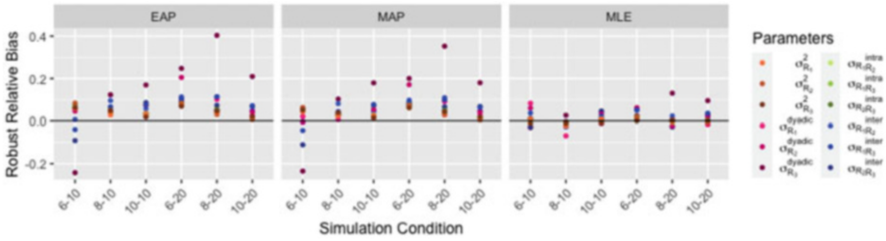

Fig. 2 Robust relative bias of point estimates at the dyad level

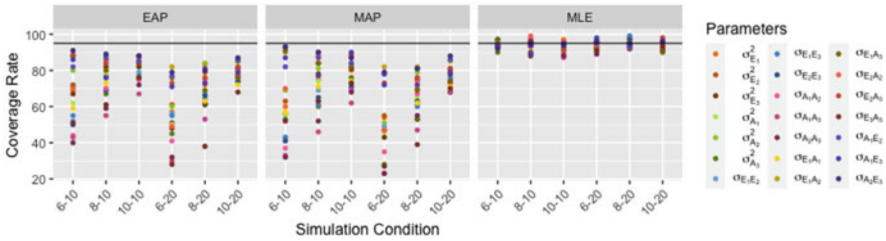

Fig. 3 Coverage rates of interval estimates at the individual level

$n_g(n_g - 1) \times G$ dyads is used to estimate relationship-level parameters. Bias for both MLE and MCMC tended to be worse in smaller groups, indicating individual-level random effects are more reliably estimated when individuals are observed across more interactions (Bonito & Kenny, 2010; Lüdtke et al., 2018).

3.2 *Accuracy of Interval Estimates*

CR for ML estimates is close to the nominal 95%, even in the small group ($n_g = 6$) conditions. CR for EAPs and MAPs at the individual level are extremely poor (see Fig. 3). Although CR for MCMC improved in larger groups, they did not converge on nominal levels at the individual level, and coverage was worse when more groups were sampled. These patterns can be attributed to the greatly underestimated parameters at the individual level and the narrowing of CIs (around inaccurate point estimates) when analyzing larger samples. CR at the dyad level was substantially better (see Fig. 4), especially in larger groups.

3.3 *Relative Efficiency of Estimators*

Figures 5 and 6 first depict the relative efficiency of the MCMC estimates with respect to the ML estimates and then depict the relative efficiency of EAP versus

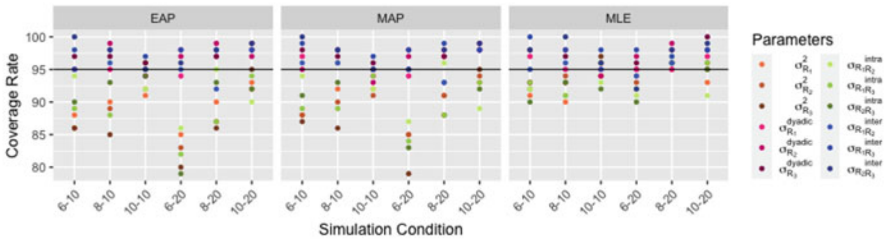

Fig. 4 Coverage rates of interval estimates at the dyad level

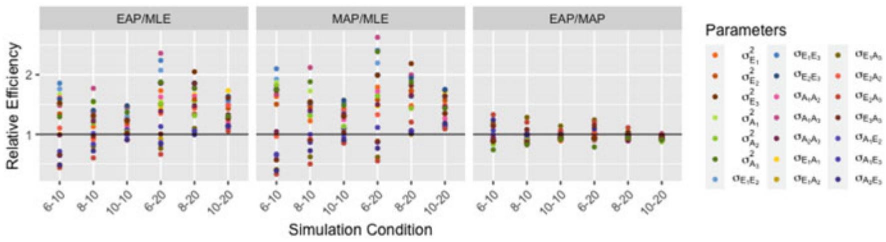

Fig. 5 Relative efficiency of point estimates at the individual level

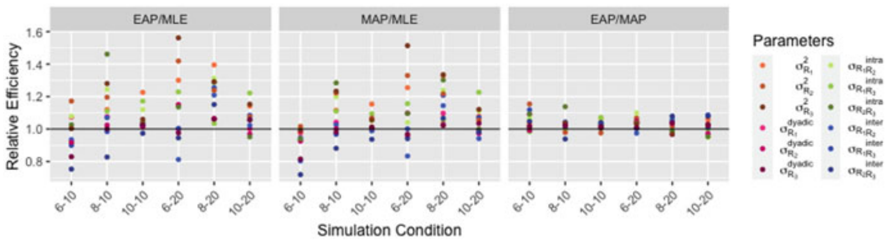

Fig. 6 Relative efficiency of point estimates at the dyad level

MAP estimates. RE $\approx$ 1 indicates that the RMSE values of both estimation approaches are approximately equal. RE $>$ 1 indicates that the estimation technique in the numerator is less efficient (has more biased estimates or displays greater sampling variability), whereas RE $<$ 1 indicates that the estimation technique in the numerator is more efficient.

At both levels, MLE tends to outperform the MCMC estimators across all conditions. EAP and MAP estimates generally perform similarly at either level. Nestler et al. (2020) computed RMSE as the square root of the sum of squares of the robust bias and median absolute deviation (MAD) of a parameter. MAD^2 provides a robust alternative to quantify the sampling variance (Talloen et al., 2019). Across all parameters at both levels, the sampling variability for MCMC and MLE is near-zero and, on average, approximately equal (see https://osf.io/w3jue/). In fact, specifically at the individual level, sampling variability for MLE is slightly higher than that of MCMC. Thus, the poor performance of MCMC relative to MLE is largely attributable to MCMC's greater magnitude of (relative) bias.

4 Discussion

We presented simulation results comparing the accuracy and efficiency of MCMC and MLE in estimating a multivariate SRM. Although both MLE and MCMC provide rather biased point estimates (particularly in small-group conditions), MCMC generally performed worse than MLE under the conditions investigated here (data are complete and normally distributed, there are no small variances, and diffuse MCMC priors are used).

One explanation for the poor performance of MCMC compared to MLE in this simulation is the combination of small sample conditions without sufficiently informative priors for the given model complexity (estimating 21 parameters at the individual level and 12 parameters at the dyad level). In their systematic review, Smid et al. (2020) found that Bayesian estimates computed based on naïve (i.e., software-default) priors led to much more bias than frequentist estimates in small samples. The diffuse priors might have influenced the posterior by placing too much weight on implausibly large values, which our small samples could not overcome.

Consistent with past research (Nestler, 2018; Lüdtke et al., 2013, 2018), we found that the accuracy of ML and MCMC estimates slightly improved as n_g increased, irrespective of G. Large groups provide more interactions per person, stabilizing individual-level estimates. However, increasing G from 10 to 20 for a given n_g did not appear to substantially improve the accuracy of estimation. This is because increasing G decreases sampling variability, giving inaccurate estimates greater precision. Hence, in a trade-off of n_g and G, fewer large groups yield more accurate estimates than many small groups (Kenny et al., 2006, p. 215).

Finally, our results indicate that the choice of EAP or MAP would not substantially affect the accuracy of (co)variance estimates when using MCMC under the conditions we investigated. This is because our population values did not contain any small or near-zero variances. When the distribution of a particular parameter is skewed (e.g., near-zero variances are a boundary condition), the EAP of the posterior provides a more extreme estimate than the MAP (Lüdtke et al., 2013). Thus, when variances are expected to be near-zero (e.g., when dyadic behavior is driven predominantly by ego or by alter effects), EAP and MAP estimates will diverge unless more informative priors are specified (Ten Hove et al., 2020).

In conclusion, small round-robin group conditions lead to biased estimates of multivariate SRM parameters when using MLE or MCMC estimation, although the latter can perform worse when using diffuse priors. More research is needed to ascertain whether (and under what conditions) MCMC estimation can provide more accurate and efficient results. One solution worth exploring is specifying more thoughtful (but weakly informative) empirical Bayes priors to stabilize the MCMC estimates.

References

Bonito, J. A., & Kenny, D. A. (2010). The measurement of reliability of social relations components from round-robin designs. *Personal Relationships, 17*(2), 235–251.

Gill, P. S., & Swartz, T. B. (2001). Statistical analyses for round robin interaction data. *Canadian Journal of Statistics, 29*(2), 321–331.

Gleason, J. R., & Halperin, S. (1975). A paired compositions model for round-robin experiments. *Psychometrika, 40*(4), 433–454.

Hoff, P. D. (2005). Bilinear mixed-effects models for dyadic data. *Journal of the American Statistical Association, 100*(469), 286–295.

Hoffman, M. D., Gelman, A., et al. (2014). The No-U-Turn Sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research, 15*(1), 1593–1623.

Kenny, D. A., Kashy, D. A., & Cook, W. L. (2006). *Dyadic data analysis*. Guilford Press.

Lewandowski, D., Kurowicka, D., & Joe, H. (2009). Generating random correlation matrices based on vines and extended onion method. *Journal of Multivariate Analysis, 100*(9), 1989–2001.

Lüdtke, O., Robitzsch, A., Kenny, D. A., & Trautwein, U. (2013). A general and flexible approach to estimating the social relations model using Bayesian methods. *Psychological Methods, 18*(1), 101.

Lüdtke, O., Robitzsch, A., & Trautwein, U. (2018). Integrating covariates into social relations models: A plausible values approach for handling measurement error in perceiver and target effects. *Multivariate Behavioral Research, 53*(1), 102–124.

Nestler, S. (2018). Likelihood estimation of the multivariate social relations model. *Journal of Educational and Behavioral Statistics, 43*(4), 387–406.

Nestler, S., Lüdtke, O., & Robitzsch, A. (2020). Maximum likelihood estimation of a social relations structural equation model. *Psychometrika, 85*(4), 870–889.

Nestler, S., Robitzsch, A., & Luedtke, O. (2022). `srm`*: Structural equation modeling for the social relations model*. R Package Version 0.4-26.

R Core Team (2023). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Salazar Kämpf, M., Liebermann, H., Kerschreiter, R., Krause, S., Nestler, S., & Schmukle, S. C. (2018). Disentangling the sources of mimicry: Social relations analyses of the link between mimicry and liking. *Psychological Science, 29*(1), 131–138.

Smid, S. C., McNeish, D., Miočević, M., & van de Schoot, R. (2020). Bayesian versus frequentist estimation for structural equation models in small sample contexts: A systematic review. *Structural Equation Modeling, 27*(1), 131–161.

Stan Development Team (2023). *RStan: The R interface to Stan*. R Package Version 2.26.23.

Talloen, W., Loeys, T., & Moerkerke, B. (2019). Consequences of unreliability of cluster means and unmeasured confounding on causal effects in multilevel mediation models. *Structural Equation Modeling, 26*(2), 191–211.

Ten Hove, D., Jorgensen, T. D., & van der Ark, L. A. (2020). Comparing hyperprior distributions to estimate variance components for interrater reliability coefficients. In *Quantitative Psychology: 84th Annual Meeting of the Psychometric Society, Santiago, Chile, 2019* (pp. 79–93). Springer.

Warner, R. M., Kenny, D. A., & Stoto, M. (1979). A new round robin analysis of variance for social interaction data. *Journal of Personality and Social Psychology, 37*(10), 1742.