Testing heterogeneity in inter-rater reliability

František Bartoš^{1*}[0000-0002-0018-5573], Patřicia Martinkov^{2,3}[0000-0003-4754-8543], and Marek Brabec³[0000-0001-6367-5791]

¹ Faculty of Arts, Charles University, Prague, Czech Republic

² Faculty of Education, Charles University, Prague, Czech Republic

³ Institute of Computer Science, Czech Academy of Sciences, Prague, Czech Republic
f.bartos96@gmail.com

Abstract. Estimating the inter-rater reliability (IRR) is important for assessing and improving the quality of ratings. In some cases, the IRR may differ between groups due to their features. To test heterogeneity in IRR, the second-order generalized estimating equations (GEE2) and linear mixed-effects models (LME) were already used. Another method capable of estimating the components for IRR is generalized additive models (GAM). This paper presents a simulation study evaluating the performance of these methods in estimating variance components and in testing heterogeneity in IRR. We consider a wide range of sample sizes and various scenarios leading to heterogenous IRR. The results show, that while the LME and GAM models perform similarly and yield reliable estimates, the GEE2 models may lead to incorrect results.

Keywords: Inter-rater reliability, Mixed-effect models, Generalized estimating equations.

1 Introduction

Ratings by multiple raters are used in assessing quality of scientific articles, grant proposals or job candidates. The credibility of ratings is contingent upon its reliability, validity, and fairness (American Educational Research Association, American Psychological Association, & National Council on Measurement in Education, 2014). One type of reliability, the inter-rater reliability (IRR), measures the degree of consistency between raters. It can be defined as the ratio of true score variance to total score variance (Lord, 1959; Novick, 1966) which in the simplest case corresponds to the intra-class correlation coefficient in a random-effect model.

Furthermore, the IRR may differ between groups. For example, Martinkov, Goldhaber, & Erosheva (2018) proposed linear mixed-effect models (LME) to account for differences in variance terms between groups and confirmed significant differences in IRR when rating internal vs. external applicants. Whereas Mutz, Bornmann, & Daniel (2012) utilized the second-order generalized estimating equations (GEE2) and confirmed differences in IRR in ratings of grant proposals from different disciplines. A question arises which of these two, or other possible approaches is superior for testing heterogeneity in IRR.

The aim of this study is to compare precision of different procedures in testing differences in IRR between two groups. We present a simulation study comparing LME, GEE and a newly considered approach based on generalized additive models (GAM). We designed a simulation study testing how do individual methods compare across scenarios in which the heterogeneity is introduced by differences in structural variances (variance of the random-effects in the LME framework), differences in residual variances, differences in means of the ratings, or their combination. Moreover, we varied the number of ratees and raters to explore how sample size influences precision of estimation of the individual model parameters and IRR itself.

2 Methods

2.1 Inter-rater reliability

In cases with nested measurements, such as the case of ratees rated by multiple raters, the IRR¹ might be estimated using a variance decomposition and calculating the intraclass correlation (ICC) (McGraw & Wong, 1996; Shrout & Fleiss, 1979). In the most trivial example, when assuming the only structural effect causing differences in ratings being the ratees themselves, the ICC can be described by a single random intercept mixed-effect model (Equation 1) with the observed j^{th} rating of i^{th} ratee Y_{ij} modeled as sum of overall mean μ , rate-specific intercept α_i and random error $\epsilon_{i,j}$

$$Y_{i,j} = \mu + \alpha_i + \epsilon_{i,j}. \quad (1)$$

The IRR is then defined as a proportion of the ratee variance σ_α^2 to the overall variance in ratings (Equation 2), which corresponds to ICC

$$IRR = \frac{\sigma_\alpha^2}{\sigma_\alpha^2 + \sigma_\epsilon^2}. \quad (2)$$

Even in this trivial case, heterogeneity in the IRR might emerge due to differences in the structural variance σ_α^2 - the variability of the ratees' true score, due to differences in the residual variance σ_ϵ^2 , or their combination. Although possible differences in the overall mean ratings μ do not affect the (2) directly, they might lead to a biased estimate if they are not accounted for in (1).

2.2 The second-order generalized estimating equations

One way of estimating the IRR is by using the GEE2. Generalized estimating equations were originally introduced as a way of dealing with clustered data by using a “working correlation matrix” easing the computation in comparison to mixed-effect models (Lipsitz & Fitzmaurice, 2008). GEE2 in comparison to the generalized estimating equations as introduced by Liang & Zeger (1986) allow not only to specify the model for the mean (Equation 3) but also the residual variance (Equation 4) and intra-class

¹ We are using IRR to refer to single-rater IRR in the paper.

correlation (Equation 5), to depend on a set of covariates (Yan & Fine, 2004). We consider only a specific case in which the covariate is group membership, thus the Equations (3), (4), (5) consist only of group-specific (g) intercepts γ transformed to appropriate parameter (group specific mean μ_g , residual variance $\sigma_{\epsilon,g}^2$ and ICC_g) by a given link function (f_1, f_2, f_3 respectively).

$$f_1(\mu_g) = \gamma_{1,g}, \quad (3)$$

$$f_2(\sigma_{\epsilon,g}^2) = \gamma_{2,g}, \quad (4)$$

$$f_3(ICC_g) = \gamma_{3,g}. \quad (5)$$

2.3 Linear mixed-effect models

LME offer a different way of accounting for clustering in observed ratings $Y_{i,j,g}$ of j^{th} rating of i^{th} individual from group g by fully specifying joint distribution within clusters via latent variables (Equation 6). In contrast to GEE2, they do not specify a model for ICC but random-effects directly, leading to estimates of the group-specific structural variances $\sigma_{\alpha,g}^2$, assuming normally distributed random effects with mean zero, group specific residual variance $\sigma_{\epsilon,g}^2$, assuming normally distributed residuals with mean zero, and group specific mean μ_g

$$Y_{i,j,g} = \mu_g + \alpha_{i,g} + \epsilon_{i,j,g}. \quad (6)$$

The mixed-effect models can be estimated either in a classical frequentist framework using maximum or restricted maximum likelihood (ML, REML) or by Markov Chain Monte Carlo (MCMC) in a Bayesian framework (Browne & Draper, 2006)

2.4 Generalized additive models

The GAM models (Wood, 2017) are generally fitted using penalized likelihood with quadratic penalties and generalized cross-validation (GCV) estimates of unknown penalization constants. LME models can be viewed as a special case of this general formulation, where the penalty matrix has (somewhat unusually in the context of GAMs motivated by smoothing) full rank, leading to more convenient computations. Penalty coefficients are related to the variance of random effects. In our context, we view (part of the) GAM framework just as a tool for alternative and flexible estimation of LME models.

3 Simulation study

3.1 Data generation

The data generating mechanism corresponds to the LME model specification in Equation 6 with a given number of ratees (I) from 2 groups (g) who are being rated J

times. Equation 6 implies the average rating of ratees (μ_g), structural variance ($\sigma_{\alpha,g}^2$) and residual variance ($\sigma_{\epsilon,g}^2$) resulting in group-specific IRR

$$IRR_g = \frac{\sigma_{\alpha,g}^2}{\sigma_{\alpha,g}^2 + \sigma_{\epsilon,g}^2}. \quad (7)$$

With values inspired by results of Martinková et al. (2018) we manipulated standardized mean differences between the groups ($\mu_2 - \mu_1 = 0, 0.4$), structural variance ratios ($\sigma_{\alpha,1}^2/\sigma_{\alpha,2}^2 = 1, 1.5$), and residual variance ratios ($\sigma_{\epsilon,1}^2/\sigma_{\epsilon,2}^2 = 1, 1.5$), while constraining the overall variance to 1 and the mean IRR across groups to 0.45. This led to 8 simulation scenarios, with scenarios 4 and 8 split into two subscenarios depending on whether the structural and residual variance ratios differed in the same or the opposite direction (Table 1). Moreover, we manipulated the number of times the ratees were rated ($J = 3, 5$) and the number of ratees per group ($I = 25, 50, 100, 200$) in each scenario. In total, 10 (scenarios including subscenarios) $\times$ 2 (number of ratings) $\times$ 4 (number of ratees) = 80 conditions were simulated, 1000 times each, implying 80,000 randomly generated datasets. Code for the data generating process is provided in the Appendix.

Table 1. The simulation scenarios setting

Scenario	μ_1	μ_2	$\sigma_{\alpha,1}$	$\sigma_{\alpha,2}$	$\sigma_{\epsilon,1}$	$\sigma_{\epsilon,2}$	IRR ₁	IRR ₂
1	0.00	0.00	0.67	0.67	0.74	0.74	0.45	0.45
2	0.00	0.00	0.67	0.67	0.67	0.82	0.50	0.40
3	0.00	0.00	0.60	0.74	0.74	0.74	0.40	0.50
4.1	0.00	0.00	0.60	0.73	0.66	0.81	0.45	0.45
4.2	0.00	0.00	0.73	0.60	0.66	0.81	0.55	0.35
5	-0.20	0.20	0.67	0.67	0.74	0.74	0.45	0.45
6	-0.20	0.20	0.67	0.67	0.67	0.82	0.50	0.40
7	-0.20	0.20	0.60	0.74	0.74	0.74	0.40	0.50
8.1	-0.20	0.20	0.60	0.73	0.66	0.81	0.45	0.45
8.2	-0.20	0.20	0.73	0.60	0.66	0.81	0.55	0.35

3.2 Model implementation

The GEE2 models were estimated in R (R Core Team, 2019) using `geepack` package (Halekoh, Højsgaard, & Yan, 2006) with fully iterated jackknife variance estimator, exchangeable covariance matrix, identity link to mean, exponential link to the residual variance, and modified Fisher-z transformation restricting the ICC to [-1, 1] interval as in Mutz et al. (2012). All the remaining settings of `geese()` function were left at their default values, with the maximal number of iterations being increased to 500 just for the case of slower convergence.

To fit LME models, we used three types of implementation as specified below. First, the `lme4` package (Bates, Mächler, Bolker, & Walker, 2015) was used as in

Martinková et al. (2018). Because `lme4` does not allow for the specification of the residual variance in LME, we also used the `nlme` package (Pinheiro, Bates, DebRoy, Sarkar, & R Core Team, 2019). Finally, Bayesian estimates through MCMC were implemented through customized models written in `Stan` (Carpenter et al., 2017) and fitted via the `rstan` package (Stan Development Team, 2019).

The LME with `lme4` package were fitted using REML with default settings. The results were bootstrapped 1000 times, as in Martinková et al. (2018), in order to obtain the standard errors for the structural and residual variance and the IRR estimates with confidence intervals (CI). In the case of non-convergence, the non-convergent fit was updated using 20,000 additional function evaluations.

In models fitted with the `nlme` package the structural variance was specified using the “pdDiag” argument in `random`, and the residual variance by the “varIdent” argument in `weights`, in order to allow variances to differ by group. The models were fitted by REML with all other settings kept at default values. All the transformations and computation required for obtaining the final estimates and their standard errors, including the one IRR, were done by a delta method implemented in the `car` package (Fox & Weisberg, 2019), which is a method for error propagation that allows to obtain an approximate distribution for a function of an asymptotically normal statistical estimator (Doob, 1935). In the case of non-convergence, the number of iterations, optimizations steps, and objective function evaluations were increased to 500 and initial estimates refinements to 50.

The Stan models were written with identity link to the dummy coded group means, structural variances and residual variances. For all parameters estimated using Stan, 95% CI were computed using the samples from the posterior distribution. After a preliminary check of computations, we used a non-central parametrization for the models with $i = 25$, and central parametrization otherwise (Betancourt & Girolami, 2015). We used a normal prior distribution with mean 0 and standard deviation 1 for the means and a half-normal prior distribution with mean 0 and standard deviation 1 for the structural and residual variances. The models were fitted using Hamiltonian Monte Carlo with two chains and iterations set to 2000 out of which 1000 was set aside for warm-up. All the remaining settings were kept at the default values. In the case of non-convergence, the average proposal acceptance probability was increased to 0.95 and the maximum treedepth to 15.

The GAM models were fitted using the `mgcv` R package (Wood, 2011) which allows estimating the random-effects using a “smoother function” – which leads to penalized likelihood estimation. The model specification for residual variance requires `gauss` family function, however, we did not manage to set the optimizer in a way that would not produce underestimated estimates for the structural variance across all our simulation scenarios. We, therefore, used model facilitating gaussian family function which does not allow to specify a model for the residual variance. The random effects were specified using a “re” type spline and the models were fitted using REML, with all the remaining settings at the default values. The IRR and its standard errors were computed using the delta method. In the case of non-convergence, the maximum number of iterations was increased to 500 Code for the models corresponding to the

most complex scenarios is provided in the Appendix, codes for all models is accessible in online Supplementary material available at <https://osf.io/9sajx/>.

3.3 Simulation evaluation

Models were deemed as non-convergent if there was a warning or error message returned while the model was fit or accessed. In the case of Stan models, we deemed model as non-convergent in following cases: any divergent iteration, an effective sample size below 100, Bayesian fraction of missing information below 0.20, or R-hat above 1.10, where R-hat is indicator of chains not converging to stationary distribution, defined as the average variance from draws of one chain to variance of draws from all chains (Gelman, Rubin, & others, 1992). Otherwise, the model was refitted using the settings specified above. The fitting times include the refitting process and the bootstrap, in case of models fitted using `lme4`, with both chains in Stan models running in parallel.

The bias of group means, structural variances, residual variances, and IRR estimates were computed as an average difference between the estimated and true values. Root mean square error (RMSE) was calculated by taking the root of the mean of squared differences between true and estimated values. Both of these indices were computed from all models corresponding to the data generating mechanism and also for only those marked as converged.

In addition, the coverage of 95% CI was computed as a proportion of CI containing the true value.

Furthermore, the methods were also compared in terms of power and error rate. In cases where the models corresponding to the data generating mechanism allowed for a particular estimate to differ between the groups, the power was computed as a proportion of significant z-tests at the $\alpha = .05$ level. In the remaining cases, the error rate for a particular estimate was computed as a proportion of significant z-tests at the $\alpha = .05$ level in a model identical to the one corresponding to the data generating process but allowing the particular parameter to vary by group². In Stan and `lme4` models, the difference in IRR was tested by an overlap of 95% CI for the difference between the by group IRR with zero.

Because there were minimal differences in all evaluation metrics between all models and only models that converged, we present only results for all models (unconditional on convergence). The differences between $J = 3$ and $J = 5$ were also rather minimal (apart from up to 10% increase in power for structural variances) therefore only results from simulations with $J = 3$ are presented. Results conditional on convergence and $J = 5$ can be found in online Supplemental Material.

² I.e., in scenario 1, the error rate for mean parameter in mixed-effect model was computed using estimates from a model corresponding to the scenario 5. Furthermore, the error rate for IRR in scenario 1 was computed only for GEE2, because only GEE2 offers the possibility to let the ICC parameter vary by group with the remaining parameters being equal across groups.

4 Results

4.1 Convergence & fitting times

All the GEE2 and GAM models converged and only 3 GEE2 models needed to be refitted. The nlme and Stan models also converged in almost all simulations. The lowest convergence for nlme models occurred in scenarios 3 and 7 for nlme models with 97.9% of models converging before and 99.7% after a refit, and Stan models not dropping below 98.8% in scenario 4.2 before and 99.4% after refitting. However, while the lme4 models almost always converged in scenario 1 and 5, with 100.0% and 99.8% convergence before refit and 100.0% after refitting, in scenario 3 and 7 the convergence dropped to 23.1% and 22.4% before and only to 39.4% and 39.3% after refitting (Supplementary Table 1). A closer look into the lme4 convergence issues revealed that a higher sample size led to worse convergence, 51.0% and 50.0% convergence after refit in scenarios 3 and 7 with $I = 25$ and $J = 3$, and 23.9% and 25.2% convergence after refit with $I = 200$ and $J = 5$, with more details in supplementary Table 2.

In regard to fitting times, the GEE2 models were fitted the fastest ($Mdn = 0.02$ sec), followed by LME models implemented in nlme ($Mdn = 0.09$ sec), the GAM models ($Mdn = 1.07$ sec), LME models implemented in Stan ($Mdn = 5.28$ sec), with the slowest models being the LME models implemented in lme4 due to bootstrapping ($Mdn = 16.42$ sec). All of the fitting times increased with sample size (see Fig. 1).

Fig. 1. Median fitting times (in seconds) of different algorithms for number of ratings $J = 3$, across scenarios (plots) and number of rates per group ($I = 25, 50, 100, 200$, see x-axis).

4.2 Bias & RMSE

The bias of the mean estimates was very low across all scenarios and methods ($< |0.009|$). However, the GEE2 produced a significant amount of a positive bias for residual standard deviations (> 0.217) and for IRR estimates (> 0.238) across all sample sizes and scenarios. The remaining methods provided reasonable estimates with a considerably lower bias for structural standard deviations ($< |0.014|$), residual standard deviations ($< |0.020|$), and IRR ($< |0.034|$) (Fig 2).

The RMSE tells a very similar story, with very low values for mean estimates across all models with a visible improvement with the sample size. As in the bias case,

the RMSE for residual standard deviation (> 0.239) and IRR (> 0.277) estimates from GEE2 models were rather high and mostly result of bias. The remaining methods produced RMSE considerably smaller in structural standard deviation (< 0.167), residual standard deviation (< 0.089) and IRR (< 0.124), with all RMSE decreasing with sample size (Fig 3).

4.3 CI coverage

The CI coverage of the mean estimates was close to the nominal bound across all models (94.3%). However, the mean CI coverage across all scenarios of GEE2 models was low not only for residual variances ($M = 2.9\%$) but also for IRR estimates ($M = 26.2\%$). The other methods performed much better, with the only CI coverage lower than 90% for one of the structural variances in scenario 4.2 and 8.2 in nlme models, with mean CI coverage of 94.7% for structural variances, 94.9% for residual variances and 94.0% for IRR across all models and sample sizes (Fig. 4).

4.4 Power & error-rate

The error rate for testing differences between group means was around the nominal level for all scenarios and on average 5.4% and power was swiftly improving with sample size across all models (Figure 5). The models fitted with GEE2 showed a high error rate (up to 48.3%) in testing the group differences in the residual standard deviations in scenarios when the structural standard deviations differed as well, which increased with sample size (scenarios 3 and 7, Figure 5). Moreover, the GEE2 models exhibited much lower power (as low as 12.6%) to detect group differences in IRR estimates in comparison to the other methods ($> 21.7\%$) when the residual variances were different (scenarios 2, 4.2, 6, 8.2). The occasional nonmonotonic patterns in GEE2 models resulted from high bias and RMSE of the estimated parameters. The remaining methods kept adequate error rate when testing group differences in structural standard deviations ($M = 4.3\%$), residual standard deviations ($M = 5.1\%$) and IRR estimates ($M = 5.0\%$). The power of LME and GAM models increased with sample size, however, while it approached 100% for both the tests of differences between residual standard deviations (max = 99.0%) and IRR (max = 99.1%), it was rather low for tests of differences between structural standard deviations (max = 53.5%).

Fig. 2. The bias of parameters (columns) from all models corresponding to the data generating mechanism (rows) and number of rating $J = 3$ across number of rates per group ($I = 25, 50, 100, 200$, see x-axis) (bias of residual variance and IRR in GEE2 models is out of the plotting range).

Fig. 3. The RMSE of parameters (columns) from all models corresponding to the data generating mechanism (rows) and number of rating $J = 3$ across number of rates per group ($I = 25, 50, 100, 200$, see x-axis) (RMSE of residual variance and IRR in GEE2 models is out of the plotting range).

Fig. 4. The CI coverage of parameters (columns) from all models corresponding to the data generating mechanism (rows) and number of rating $J = 3$ across number of rates per group ($I = 25, 50, 100, 200$, see x-axis) (the CI coverage of residual variance and IRR in GEE2 models is out of the plotting range).

Scenario

Fig. 5. The error-rate and power for a z-test of difference between parameters (columns) from all models corresponding to the data generating mechanism (rows) and number of rating $J = 3$ across number of ratees per group ($I = 25, 50, 100, 200$, see x-axis).

5 Discussion

This study compared the GEE2, LME and GAM models in estimating IRR and other parameters. As a main result, we uncovered an unsatisfactory performance of GEE2 implemented in the `geepack` package in estimating group residual variances and IRR as measured by bias, RMSE, and CI coverage, and in testing differences in group residual variances and IRR as measured by error-rate, and power. This may be due to less precise model specification in GEE2 than in LME and GAM models. We might only speculate that some fine-tuning of `geese.control` tolerances could lead to improvements in the quality of fit (while possibly also worsening convergence rates), but in this study we insisted on using default values as a proxy for behavior of a typical user. Nevertheless, the estimates of group means produced by GEE2 and also by all other methods were precise and with both standard errors improving with sample size.

The LME implemented in the `lme4` package exhibited the worst convergence behavior, especially in cases with different structural variances by groups and higher sample sizes. The results from the non-converged models did not seem to affect the results³, hinting on a possible difference in the strictness of the convergence checks. Finally, the models implemented in `lme4` was their fitting times, due to bootstrapping when calculating CI for variance components and IRR. With the biggest drawback being their inability to estimate models with different parameters for residual variances across groups.

On the other hand, models implemented in the `nlme` package showed almost perfect convergence, allowed to fit a model with different parameters for residual variances across groups and displayed the second-fastest fitting times next to GEE2. On top of that, the estimated parameters were satisfactory and comparable to the remaining LME methods and GAM models. Nonetheless, one of the drawbacks of `nlme` might be the difficulty of obtaining CI for the IRR estimates, requiring numerous transformations and the delta method.

In comparison to models implemented in the `nlme` package, the models implemented in Stan required a longer time to fit and their coding was more challenging. However, this burden might be overcome, for example, by the `brms` package (Bürkner, 2017) and when using the Bayesian approach, the subsequent manipulation with the estimates was much simpler.

The time required to fit the GAM models was on the higher end, especially with increasing number size. But there were no problems with convergence and models yielded estimates with similar qualities, power and error rate as those obtained by LME models. Moreover, the GAM models allow formulating a wider range of models than LME models, thus creating further possibilities for future applications.

Furthermore, it is important to note that a small number of ratings for each ratee implied a low power to detect differences in the structural variances between the groups no matter which method was used. An increasing number of ratings per ratee led to a

³ The authors do NOT advise to ignore convergence warnings. These results might be solely due to the specific simulation scenarios settings and models fitted.

noticeable improvement, however, more than 5 ratings per ratee would be still needed to achieve adequate power.

There are several limitations worth mentioning. In order to make the simulations feasible, the study relied on the most simplified data generating processes possible. This way, we were able to assess the difficulty of estimating the individual components and their combinations, however, the generalizability of the findings to more complex data scenarios is unsure. Further studies will be needed to explore more complicated scenarios. Furthermore, throughout the study, it was assumed that the data generating process is known and the only issue to solve is a proper estimation of parameters of the given model or testing differences in given parameters. That is scarcely the case in the most real-life problems, where a substantial theory and/or model selection mechanisms are needed. While from ordinary user perspective testing significance of group effects seems straightforward in GEE2 even in case of more effects or their interactions, model-building and testing may seem more complicated in LME and GAM framework. All that said, the presented simulations may serve as a helpful guideline for estimating and testing differences in IRR between groups.

6 Conclusions

This simulation study has shown that while both LME and GAM are reliable methods for estimating differences in IRR between groups, GEE2 can lead to biased results, inadequate coverage of CI and tests with high error-rates. Therefore, we advise that either LME or GAM is used for testing heterogeneity in IRR.

Acknowledgment

The work was partially supported by grants PRIMUS/17/HUM/11 and SVV 2019 – 260482 realized at the Charles University, Faculty of Education and Faculty of Arts and by the long-term strategic development financing of the Institute of Computer Science (Czech Republic RVO 67985807). Access to computing and storage facilities owned by parties and projects contributing to the National Grid Infrastructure MetaCentrum provided under the programme "Projects of Large Research, Development, and Innovations Infrastructures" (CESNET LM2015042), is greatly appreciated.

Appendix

Function for data simulation takes vectors of length two as an argument for group-specific means (`mu`), structural standard deviation (`sigma_alpha`), residual standard deviation (`sigma_epsilon`), and two integers (`I`, `J`) defining the number of ratees in each group `I` and number of ratings per ratee `J`.

```
simulate_data <- function(mu, sigma_alpha, sigma_epsilon,
```

```

I, J){
  # create group membership
  g <- c(rep(1,I), rep(2,I))
  # random effects / ratees' true scores
  alpha <- rnorm(2*I,mu[g],sigma_alpha[g])
  # ratees' indices
  id <- rep(1:(2*I),J)
  # observed values Y
  Y<-rnorm(length(id),alpha[id],sigma_epsilon[g[id]])
  data <- cbind.data.frame("id"=id,"group"=g,"y"=Y)

  data <- data[order(data$id),]
  return(data)
}

```

Thus, the data for the first simulation scenario with 25 ratees rated 3 times could be generated as follows:

```

data <- simulate_data(mu = c(0.00, 0.00),
                      sigma_alpha = c(0.67, 0.67),
                      sigma_epsilon = c(0.74, 0.74),
                      I = 25, J = 3)

```

Functions for fitting models corresponding to the most complex scenario 8:

```

### GEE2 model corresponding to scenario 8
# design matrix for correlation parameters
z <- with(data[!duplicated(data$id),],
          model.matrix(~
as.factor(group) - 1))
gee_model <- geepack::geese(y ~ as.factor(group) - 1,
                          zcor = z, corstr = "exchangeable",
                          sformula = ~ as.factor(group) - 1,
                          id = id, data = data,
                          mean.link = "identity",
                          sca.link = "log",
                          cor.link = "fisherz",
                          family = gaussian)

### nlme model corresponding to scenario 8
lme_model <- nlme::lme(y ~ as.factor(group) - 1,
                      random= list(id = nlme::pdDiag(
form = ~ as.factor(group) - 1)),
                      weights = nlme::varIdent(
form = ~ 1 | as.factor(group)),

```

```

data = data, method = "REML")

### lme4 model corresponding to scenario 4
lmer_model <- lme4::lmer(y ~ as.factor(group) - 1
                        + (0 + as.factor(group)|id),
                        data = data, REML = TRUE)

### Stan model corresponding to scenario 8
# stan code with non-central parametrization
stan_code_l11_nc <- c('
data {
    int<lower=0> N;
    int<lower=0> N_id;
    int id[N];
    int group[N];
    int id_group[N_id];
    vector[N] y;
}
parameters {
    vector[2] mu_group;
    vector[N_id] alpha_z;
    vector<lower=0>[2] sigma_alpha;
    vector<lower=0>[2] sigma_epsilon;
}
transformed parameters {
    vector[N_id] alpha;

    for(j in 1:N_id){
        alpha[j] = alpha_z[j]*sigma_alpha[id_group[j]];
    }
}
model {
    vector[N] mu;
    vector[N] sigma;

    target += normal_lpdf(mu_group | 0, 1);
    target += normal_lpdf(sigma_alpha | 0, 1)
              - normal_lccdf(0 | 0, 1);

    target += normal_lpdf(sigma_epsilon | 0, 1)
              - normal_lccdf(0 | 0, 1);
    target += normal_lpdf(alpha_z | 0, 1);

    for(n in 1:N){
        mu[n] = mu_group[group[n]] + alpha[id[n]];
    }
}

```



```

        sigma[n] = sigma_epsilon[group[n]];
    }

    target += normal_lpdf(y | mu, sigma);
}
')
# formatting data for stan
data_stan <- list(
  "id"      = data$id,
  "group"   = as.integer(data$group),
  "y"       = data$y,
  "id_group" = data$group[!duplicated(data$id)],
  "N"       = nrow(data),
  "N_id"    = length(unique(data$id))
)
# compiling model
stan_111_nc <- rstan::stan_model(model_code =
                                stan_code_111_nc)
# sampling the model
model_stan <- rstan::sampling(stan_111_nc,
                              data = data_stan,
                              iter = 2000, warmup = 1000,
                              chains = 2, cores = 2)

# GAM model corresponding to data generating scenario 8
# remapping the ids for more efficient fitting
data$id_new <- as.factor(ifelse(data$group == 1, data$id,
                                data$id - length(unique(data$id))))
data$group <- as.factor(data$group)
mgcv_model <- mgcv::gam(y ~ group - 1
                      + s(id_new, bs = "re", by = group),
                      data = data, method = "REML")

```

References

- American Educational Research Association, American Psychological Association, & National Council on Measurement in Education. (2014). *Standards for Educational and Psychological Testing*. Washington, DC: American Educational Research Association.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67(1), 1–48.
<https://doi.org/10.18637/jss.v067.i01>
- Betancourt, M., & Girolami, M. (2015). Hamiltonian Monte Carlo for hierarchical models. *Current Trends in Bayesian Methodology with Applications*, 79, 30.

- Browne, W. J., & Draper, D. (2006). A comparison of Bayesian and likelihood-based methods for fitting multilevel models. *Bayesian Analysis*, 1(3), 473–514. <https://doi.org/10.1214/06-BA117>
- Bürkner, P.-C. (2017). brms: An R Package for Bayesian Multilevel Models Using Stan. *Journal of Statistical Software*, 80(1), 1–28. <https://doi.org/10.18637/jss.v080.i01>
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., ... Riddell, A. (2017). Stan: A Probabilistic Programming Language. *Journal of Statistical Software*, 76(1), 1–32. <https://doi.org/10.18637/jss.v076.i01>
- Doob, J. L. (1935). The Limiting Distributions of Certain Statistics. *The Annals of Mathematical Statistics*, 6(3), 160–169. <https://doi.org/10.1214/aoms/1177732594>
- Fox, J., & Weisberg, S. (2019). *An R Companion to Applied Regression* (Third). Retrieved from <https://socialsciences.mcmaster.ca/jfox/Books/Companion/>
- Gelman, A., Rubin, D. B., & others. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, 7(4), 457–472.
- Halekoh, U., Højsgaard, S., & Yan, J. (2006). The R package geepack for generalized estimating equations. *Journal of Statistical Software*, 15(2), 1–11.
- Liang, K.-Y., & Zeger, S. L. (1986). Longitudinal data analysis using generalized linear models. *Biometrika*, 73(1), 13–22. <https://doi.org/10.1093/biomet/73.1.13>
- Lipsitz, S., & Fitzmaurice, G. (2008). Generalized estimation equations for longitudinal data analysis. *Longitudinal Data Analysis*. New York: Chapman & Hall/CRC, 43–78.
- Lord, F. M. (1959). Statistical inferences about true scores. *Psychometrika*, 24(1), 1–17. <https://doi.org/10.1007/BF02289759>
- Martinková, P., Goldhaber, D., & Erosheva, E. (2018). Disparities in ratings of internal and external applicants: A case for model-based inter-rater reliability. *PloS One*, 13(10), e0203002. <https://doi.org/10.1371/journal.pone.0203002>
- McGraw, K. O., & Wong, S. P. (1996). Forming inferences about some intraclass correlation coefficients. *Psychological Methods*, 1(1), 30–46. <https://doi.org/10.1037/1082-989X.1.1.30>
- Mutz, R., Bornmann, L., & Daniel, H.-D. (2012). Heterogeneity of Inter-Rater Reliabilities of Grant Peer Reviews and Its Determinants: A General Estimating Equations Approach. *PLOS ONE*, 7(10), e48509. <https://doi.org/10.1371/journal.pone.0048509>
- Novick, M. R. (1966). The axioms and principal results of classical test theory. *Journal of Mathematical Psychology*, 3(1), 1–18. [https://doi.org/10.1016/0022-2496\(66\)90002-2](https://doi.org/10.1016/0022-2496(66)90002-2)
- Pinheiro, J., Bates, D., DebRoy, S., Sarkar, D., & R Core Team. (2019). *nlme: Linear and Nonlinear Mixed Effects Models*. Retrieved from <https://CRAN.R-project.org/package=nlme>
- R Core Team. (2019). *R: A Language and Environment for Statistical Computing*. Retrieved from <https://www.R-project.org/>
- Shrout, P. E., & Fleiss, J. L. (1979). Intraclass correlations: Uses in assessing rater reliability. *Psychological Bulletin*, 86(2), 420–428.

- Stan Development Team. (2019). *RStan: The R interface to Stan*. Retrieved from <http://mc-stan.org/>
- Wood, S. N. (2011). Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society (B)*, 73(1), 3–36.
- Wood, Simon N. (2017). *Generalized additive models: An introduction with R*. Chapman and Hall/CRC.
- Yan, J., & Fine, J. (2004). Estimating equations for association structures. *Statistics in Medicine*, 23(6), 859–874. <https://doi.org/10.1002/sim.1650>