PSYCHOMETRIKA—VOL. 85, NO. 3, 815–836
SEPTEMBER 2020
https://doi.org/10.1007/s11336-020-09718-1

A DYADIC IRT MODEL

BRIAN GIN

UNIVERSITY OF CALIFORNIA, SAN FRANCISCO

NICHOLAS SIM

UNIVERSITY OF CALIFORNIA, BERKELEY

ANDERS SKRONDAL

NORWEGIAN INSTITUTE OF PUBLIC HEALTH

UNIVERSITY OF OSLO

UNIVERSITY OF CALIFORNIA, BERKELEY

SOPHIA RABE- HESKETH

UNIVERSITY OF CALIFORNIA, BERKELEY

We propose a dyadic Item Response Theory (dIRT) model for measuring interactions of pairs of individuals when the responses to items represent the actions (or behaviors, perceptions, etc.) of each individual (actor) made within the context of a dyad formed with another individual (partner). Examples of its use include the assessment of collaborative problem solving or the evaluation of intra-team dynamics. The dIRT model generalizes both Item Response Theory models for measurement and the Social Relations Model for dyadic data. The responses of an actor when paired with a partner are modeled as a function of not only the actor's inclination to act and the partner's tendency to elicit that action, but also the unique relationship of the pair, represented by two directional, possibly correlated, interaction latent variables. Generalizations are discussed, such as accommodating triads or larger groups. Estimation is performed using Markov-chain Monte Carlo implemented in `Stan`, making it straightforward to extend the dIRT model in various ways. Specifically, we show how the basic dIRT model can be extended to accommodate latent regressions, multilevel settings with cluster-level random effects, as well as joint modeling of dyadic data and a distal outcome. A simulation study demonstrates that estimation performs well. We apply our proposed approach to speed-dating data and find new evidence of pairwise interactions between participants, describing a mutual attraction that is inadequately characterized by individual properties alone.

Key words: item response theory, social relations model, dyadic data, Markov-chain Monte Carlo, `Stan`.

1. Introduction

How individuals interact within a group has been and continues to be of interest to researchers in the behavioral sciences. A model developed to handle the simplest case of two individuals interacting in a dyad is the Social Relations Model (SRM) (e.g., Warner et al. 1979; Kenny and La Voie 1984), also known as the complete SRM to distinguish it from a popular constrained version called the basic SRM (see Sect. 2.4 on variants of the SRM). Here, the ways one individual

Electronic supplementary material The online version of this article (https://doi.org/10.1007/s11336-020-09718-1) contains supplementary material, which is available to authorized users.

Correspondence should be made to Nicholas Sim, University of California, Berkeley, 2121 Berkeley Way, Berkeley, CA94720, USA. Email: sim_nicholas@berkeley.edu

(often called an actor or perceiver) of a dyad behaves when paired with the other (often called the partner or target) and vice versa are analyzed to infer individual-level and dyad-level effects. The behavior of the actor can be directed toward the partner (e.g., an individual's perception of the partner's attractiveness) or undirected (e.g., the number of times an individual takes the lead in a collaborative problem-solving task) and can be measured during or after socially interacting with the partner. Compared with traditional "isolated" models where only the actor contributes to the eventual behavior, the innovative SRM includes contributions from both actor and partner, as well as their interaction, to the actor's behavior. The SRM model has been most often used in social psychology (e.g., Kenny and Kashy 1994), but is increasingly being used in other fields. A diverse set of examples include relationships in pharmacy and therapeutics hospital-committee decision-making (Bagozzi and Ascione 2005), social media ties among basketball teammates (Koster and Brandy 2018), and militarized interstate disputes (Dorff and Ward 2013).

In the original formulation of the complete SRM, the specific behavior of an actor when paired with a partner depends on a composite dyad-level latent variable composed of three parts: (i) an individual-level latent trait reflecting a general inclination of the actor to behave in a certain way when paired with a partner, (ii) an individual-level latent trait reflecting the general tendency of the partner to elicit such a behavior, and (iii) a dyad-level latent trait that characterizes the effect of the unique (directed) relationship between both parties on the behavior of the actor that is independent of the two individual-level latent traits (Back and Kenny 2010). More concretely, if one is interested in the level of physical attraction of an actor toward a partner, then the three components reflect (i) how, on average, an actor tends to find others attractive, (ii) how, on average, the partner tends to be found attractive, and (iii) how the actor uniquely finds the partner attractive. As a result of this formulation, the SRM is identifiable only if individuals belong to multiple pairs.

The dyadic IRT (dIRT) model introduced in this paper extends the SRM to the situation where the behavior or perception of the actor is a latent variable measured by multiple indicators, items, or measures. To our knowledge, this is a new SRM for multiple items. The most common extension of the SRM to handle multiple items (or response variables) is the multivariate SRM (e.g., Kenny 1994; Card et al. 2008; Nestler 2018). This model effectively corresponds to a set of univariate basic SRMs with additional correlations of individual-level and dyad-level latent traits across items. When there are more than two or three items, the multivariate SRM has an abundance of cross-variable correlations that are challenging to interpret. Moreover, multivariate SRMs do not provide a means for predicting or scoring actor, partner and dyadic effects on an underlying latent trait. The SRM with stable construct effects (Kenny 1994; Bonito and Kenny 2010) comes closer to the dIRT but differs from it by (1) including "unstable" item-specific latent variables for the actor, partner, and dyadic effects, and (2) considering continuous responses only and assuming that the measurement model is a classical test theory model, whereas any IRT model can be used in dIRT.

IRT is the standard approach for modeling the relationship between the latent traits of individuals and their responses to a set of items in psychometrics. There are a variety of IRT models that may differ, among other things, in terms of the numbers of parameters in the model, the type of link function used, or the approach taken (e.g., confirmatory or exploratory) (see, e.g., van der Linden 2016). However, existing models treat the latent trait as a property of the individuals who responded to the items, and perhaps an external party like a rater, but do not include a unique interaction between individuals in a dyad.

Although traditional SRM and IRT models each have limitations that could be overcome by the other, there is, to our knowledge, no prior work on integrating the models. Only two related cases appear to exists: Alexandrowicz (2015) extended the Actor–Partner Interdependence Model (APIM) by Kenny (1996) and Common Fate Model (CFM) of Kenny and La Voie (1984) to work within an IRT framework. While these models relax the condition that only an individual's latent ability affects the individual's response to an item, neither of them models the dyadic interaction as a latent trait of the dyad. Furthermore, the APIM and CFM are confined to a dyadic design where each individual is paired with only one other partner, whereas the SRM handles the case where individuals belong to multiple pairs (Kenny et al. 2006).

Our contributions include the following. First, we describe our proposed dIRT model that incorporates the key features of both the complete SRM and IRT. The model includes individual- and dyad-level latent traits and corresponding variance and covariance parameters afforded by the complete SRM, while retaining all the important measurement properties of IRT. We also indicate how the model can be extended to larger groupings than dyads, such as triads. Second, we provide a literature review of related classes of models and discuss data designs and conditions for identifiability. Importantly, unlike the complete SRM, the dIRT model is identified for cross-sectional data. Third, we extend the dIRT model to let the latent traits affect a distal outcome and depend on observed covariates and cluster-level random effects. Finally, we demonstrate the practical utility of the model by applying it to a speed dating dataset and making `Stan` code available, together with a case study explaining the code. While univariate basic SRMs for one Likert scale item at a time, treated as continuous, have been applied to speed-dating data (e.g., Ackerman et al. 2015), our multivariate model accommodates the ordinal nature of the responses and allows estimation of the unique interaction variance separate from the error variance. We hope that our contributions will inspire researchers to collect and analyze dyadic data in new settings.

The structure of the paper is as follows. In Sect. 2, we introduce the basic dIRT model, discuss data design and identification, propose various extensions of the basic model, and provide a review of related models. We present a Markov-chain Monte Carlo approach to estimate the model in Sect. 3, using `Stan`, which can be run from `R` via the `rstan` function. In Sect. 4, we apply our model and estimation method to a publicly available speed-dating dataset. In Sect. 5, we conduct a simulation study to evaluate the performance of our estimator under a variety of conditions. Finally, we make some concluding remarks in Sect. 6.

2. Dyadic Item Response Theory (dIRT)

2.1. Basic dIRT Model

In a social setting where N individuals interact in groups of size n, it is likely that the behavior of individual $a \in \{1, 2, \dots, n\}$ (called the actor) in group g is affected by not only his/her own latent traits, but also those of the individuals he/she interacts with. Additionally, there could also be a "unique" component attributable to the specific composition of the group that could affect the actor's behavior above and beyond the effects at the individual level. We can extend any IRT model to deal with such a setting by replacing the latent trait θ_a of individual a with a composite latent variable $\theta_{a,g}$ of individual a in the context of group g of size n:

$$\theta_{a,g} \equiv \alpha_a + \sum_{\substack{j=1 \\ j \neq a}}^{n} \beta_j + \sum_{k \in K} \gamma_{a,g(k)}. \tag{1}$$

Here, α_a represents the inclination of the actor to behave in a certain way, β_j represents the tendency of another member j of the group to elicit the behavior, and $\gamma_{a,g(k)}$ represents the unique way members of subgroup $g(k)$ interacted to elicit the behavior from actor a. The last sum above is over all possible subgroups $g(k)$ of sizes 1 to $n-1$ excluding the actor. (The index set is defined as $K := \{A \subseteq \{1, 2, \dots, n\} \backslash \{a\} \mid |A| \geq 1\}$, i.e., the set of all subsets of $\{1, 2, \dots, n\} \backslash \{a\}$ except the empty set). Note that $\gamma_{a,g(k)}$ may include physical interactions between actor a and the other members of the group or how the behavior of actor a is altered by the mere presence of the rest of the group. For example, in a collaborative problem-solving task, α_a could represent the inclination of actor a to be vocal, β_j how much partner j tends to elicit opinions from actors, and $\gamma_{a,g(k)}$ how vocal the actor is due to the composition of the group $g(k)$. In practice, it may not be necessary to include group sizes $|A|$ larger than 2, i.e., to consider anything more than pairwise and possibly three-way interactions.

To simplify notation, in the rest of the paper, we focus on the case when $n = 2$ as it is clear how the model can be extended when working with larger group sizes. In this dyadic setting, for actor a and partner p, the composite latent variable is modeled as

$$\theta_{a,p} \equiv \alpha_a + \beta_p + \gamma_{a,p}.$$

Unlike (1) where the composite latent variable θ and in particular the dyad-level latent trait γ are indexed by the actor and the group, we can instead index θ and γ by both individuals a and p since the index set K reduces to the singleton set $\{\{p\}\}$. Here, α_a is the actor latent trait (sometimes called actor effect), β_p the partner latent trait (sometimes called partner effect), and $\gamma_{a,p}$ the dyadic latent trait (sometimes called interaction or relationship effect) which represents the unique contribution of pairing actor a with partner p to the behavior of the actor. Note that $\gamma_{a,\,p}$ is not assumed to be identical to $\gamma_{p,\,a}$ when the roles of actor and partner are reversed.

We can use any traditional IRT model for measuring $\theta_{a,p}$. The model for response $y_{a,p,i}$ to item i by actor a, when paired with the partner p, is of the form

$$g(\mathbb{P}[y_{a,p,i} = j \mid \theta_{a,p}, \boldsymbol{\xi}_{i,j}]) = f(\theta_{a,p}, \boldsymbol{\xi}_{i,j})$$

for some link function $g(\cdot)$, item parameters $\boldsymbol{\xi}_{i,j}$ (possibly including discrimination parameters or factor loadings), and functional form $f(\cdot)$. In this article, we make the local independence assumption that the item responses are conditionally independent given the composite latent variable.

For ordinal responses, we can consider the standard partial credit model (Masters 1982) by using the adjacent-category logit link, letting $\boldsymbol{\xi}_{i,j}$ represent (unidimensional) step difficulty parameters, and taking $f(\cdot)$ to be the identity function. If item i has m_i categories (from 0 to $m_i - 1$), the model becomes

$$\log\left(\frac{\mathbb{P}_{\text{PCM}}(y_{a,p,i} = j \mid \theta_{a,p}, \delta_{i,j})}{\mathbb{P}_{\text{PCM}}(y_{a,p,i} = j-1 \mid \theta_{a,p}, \delta_{i,j})}\right) = \theta_{a,p} - \delta_{i,j} \equiv (\alpha_a + \beta_p + \gamma_{a,p}) - \delta_{i,j}, \tag{2}$$

subject to the constraint that $\sum_{j=0}^{m_i-1} \mathbb{P}_{\text{PCM}}(y_{a,p,i} = j \mid \theta_{a,p}, \delta_{i,j}) = 1$, where $j \in \{1, 2, \ldots, m_i - 1\}$ and $\delta_{i,j}$ are item step difficulties. Note that we condition on $\delta_{i,j}$ because we will adopt a (pragmatic) Bayesian perspective (see Sect. 3).

In the dIRT model, the actor and partner latent traits are bivariate normal,

$$\begin{bmatrix} \alpha_a \\ \beta_a \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mu_\alpha \\ \mu_\beta \end{bmatrix}, \begin{bmatrix} \sigma_\alpha^2 & \rho_{\alpha\beta}\sigma_\alpha\sigma_\beta \\ \rho_{\alpha\beta}\sigma_\alpha\sigma_\beta & \sigma_\beta^2 \end{bmatrix}\right), \tag{3}$$

independent from bivariate normal dyadic latent traits,

$$\begin{bmatrix} \gamma_{a,p} \\ \gamma_{p,a} \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mu_\gamma \\ \mu_\gamma \end{bmatrix}, \begin{bmatrix} \sigma_\gamma^2 & \rho_\gamma\sigma_\gamma^2 \\ \rho_\gamma\sigma_\gamma^2 & \sigma_\gamma^2 \end{bmatrix}\right). \tag{4}$$

The parameters are (i) the variances σ_α^2, σ_β^2, and σ_γ^2 of the individual and dyad latent traits, (ii) the expectations μ_α, μ_β, and μ_γ of each of the individual and dyadic latent traits, and (iii) the correlations $\rho_{\alpha\beta}$ and ρ_γ. The individual-level correlation $\rho_{\alpha\beta}$ (sometimes called the general or individual reciprocity) relates the tendency of an individual to behave in a certain way (i.e., α_a or α_p) to that same individual's tendency to elicit the behavior from his/her partner (i.e., β_a or β_p). The dyad-level correlation ρ_γ (sometimes called dyadic reciprocity) relates the two (directed) latent traits of each dyad (i.e., $\gamma_{a,p}$ and $\gamma_{p,a}$) to each other.

We will extend the dIRT model in Sect. 2.3 after discussing data design and identification issues that will motivate and justify some of the extensions. The model for the latent variables in (4) assumes that the dyads are exchangeable such as same gender twins, and Sect. 2.3 discusses how to modify the model when they are not, for instance when we want to distinguish between mother–father and mother–child dyads.

2.2. *Data Design and Identification*

The dIRT model has five variance and correlation parameters (σ_α^2, σ_β^2, σ_γ^2, $\rho_{\alpha\beta}$, ρ_γ) for the individual and dyadic latent traits that imply five unique variances and covariances for the composite latent variable: one constant variance, $\mathrm{var}(\theta_{a,p}) = \sigma_\alpha^2 + \sigma_\beta^2 + \sigma_\gamma^2$, and four distinct nonzero covariances, $\mathrm{cov}(\theta_{a,p}, \theta_{p,a}) = 2\rho_{\alpha\beta}\sigma_\alpha\sigma_\beta + \rho_\gamma\sigma_\gamma^2$, $\mathrm{cov}(\theta_{a,p}, \theta_{a,q}) = \sigma_\alpha^2$, $\mathrm{cov}(\theta_{a,p}, \theta_{b,p}) = \sigma_\beta^2$, and $\mathrm{cov}(\theta_{a,p}, \theta_{b,a}) = \rho_{\alpha\beta}\sigma_\alpha\sigma_\beta$ (where a, p, b, q are all different individuals). It is straightforward to find unique solutions for the five model parameters from the five equations above, showing that the parameters are identified if the covariance matrix of the composite latent variable is identified.

This covariance matrix is identified if all the pairs of dyads involved in the covariances exist, i.e., actor/partner role reversal (sometimes referred to as "reciprocals") must occur to identify $\mathrm{cov}(\theta_{a,p}, \theta_{p,a})$, and it must be possible to belong to more than one dyad. Specifically, it must be possible for actors to be paired with several partners to identify $\mathrm{cov}(\theta_{a,p}, \theta_{a,q})$, for partners to be paired with several actors to identify $\mathrm{cov}(\theta_{a,p}, \theta_{b,p})$, and for an actor paired with a partner p to also occur in a dyad as a partner of an actor $b \neq p$ to identify $\mathrm{cov}(\theta_{a,p}, \theta_{b,a})$. It is necessary to set some mean parameters and/or step difficulty parameters to constants for identification. Here, we set the expectations of the latent traits to zero ($\mu_\alpha = \mu_\beta = \mu_\gamma = 0$) and allow the item step difficulties $\delta_{i,j}$ to be unconstrained (anchoring on latent trait scores instead of item difficulties), except that $\delta_{i,0} = 0$.

We have assumed that dyads and individuals within dyads are exchangeable by restricting the mean of $\theta_{a,p}$ to be constant (set to 0 to identify the step difficulties), by assuming that the variance of the dyadic composite latent variables is also constant, and by letting the covariances for pairs of dyads (a, p) and (b, q) depend only on the actor/partner roles of the individuals (or individual) that are present in both dyads. The corresponding five parameters σ_α^2, σ_β^2, σ_γ^2, $\rho_{\alpha\beta}$, and ρ_γ enforce no other constraints besides exchangeability and positive semi-definiteness. Li and Loken (2002) make the point, for a traditional SRM, that the model is in that sense justified by exchangeability.

When different individuals in the dyads have different roles (e.g., father and child) or when interest centers on asymmetric relationships (e.g., supervisor and trainee), the exchangeability restrictions enforced by the model are no longer justified, and we discuss how to relax them in Sect. 2.3.1. A special case of non-exchangeability is where each dyad is composed of two types of individuals, such as husbands and wives, and these types are the same across dyads, so that there cannot, for example, be husband and wife dyads as well as father and daughter dyads. Kenny et al. (2006) refer to this design as distinguishable dyads.

We now explore several dyadic designs for which the SRM is identified, following Kenny and La Voie (1984) and Malloy and Kenny (1986). The simplest and most common design is the round-robin design. In this design, each individual belongs to a dyad with every other member of

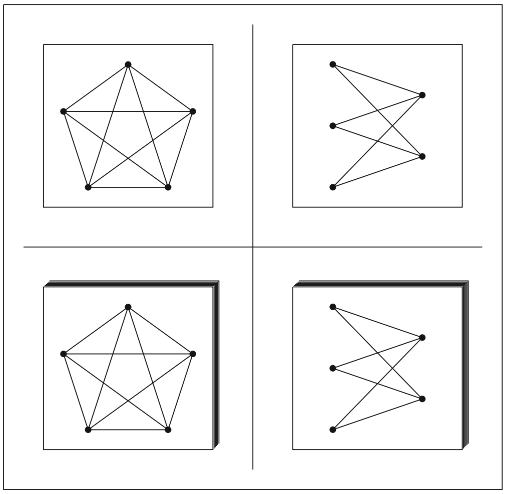

FIGURE 1.
Graphs representing round-robin design (upper-left panel), block design (upper-right panel), k-group round-robin design (lower-left panel), and k-group block design (lower-right panel). For the k-group designs, each group is represented by a layer.

the study, and there are a total of $\frac{N(N-1)}{2}$ dyads and $N(N-1)$ directed dyads. In graph-theoretic language where we view each individual as a node, the round-robin design is represented by a complete graph in the undirected case (see upper-left panel of Fig. 1) and a complete directed graph (digraph) in the directed case.

One immediate extension of the round-robin design is the block design where the N individuals are split into two blocks of sizes p and q, respectively, and $p+q=N$. Then, each individual from one block forms a dyad with every individual from the other block, but not with individuals in his/her block. That is, there are a total of pq undirected dyads and $2pq$ directed dyads. In graph-theoretic terms, such a design can be represented by a complete bipartite graph (see upper-right panel of Fig. 1). This occurs most naturally for distinguishable dyads, for example when each dyad includes one male and one female. In this case, the N individuals are split into two blocks by their gender. Kenny et al. (2006) refer to such a design as an asymmetric block design.

When individuals are nested in groups, such as families, work groups, or social networks, where each individual from the group forms a dyad with each other individual of the group, we have a "k-group round-robin design" (see lower-left panel of Fig. 1). In addition to such naturally occurring groups, the groups can also be created by the researcher to reduce response burden and costs by reducing the number of partners per actor and the number of dyads, respectively. Another reason for creating groups artificially is to allow individuals to interact within a group to create the context for the dyadic responses. For example, Christensen and Kashy (1998) created an initial social situation for groups of four lonely individuals that involved problem-solving tasks and subsequently collected dyadic ratings on personal characteristics. There can also be a block design within each group, resulting in the "k-group block design" (see lower-right panel of Fig. 1). This is the data design for the speed-dating application in Sect. 4.

2.3. *Extended dIRT Model*

2.3.1. Including Covariates for the Latent Traits The dIRT model can be extended to take into account individual and dyadic covariates that may affect the latent traits at both the individual and dyadic levels by generalizing the idea of explanatory item response models (e.g., De Boeck and Wilson 2004). One way that this can be accomplished is by specifying how the means μ_α, μ_β, and μ_γ depend on covariates, such as

$$\mu_{\alpha,a} = \boldsymbol{x}'_{\alpha,a}\boldsymbol{c}_\alpha, \quad \mu_{\beta,p} = \boldsymbol{x}'_{\beta,p}\boldsymbol{c}_\beta, \quad \mu_{\gamma,a,p} = \boldsymbol{x}'_{\gamma,a,p}\boldsymbol{c}_\gamma, \tag{5}$$

where $\boldsymbol{x}_{\alpha,a}$ are the covariates for α_a, $\boldsymbol{x}_{\beta,p}$ are the covariates for β_p, $\boldsymbol{x}_{\gamma,a,p}$ are the covariates for $\gamma_{a,p}$, and $\boldsymbol{c}_\alpha$, $\boldsymbol{c}_\beta$ and $\boldsymbol{c}_\gamma$ are the corresponding regression coefficients.

If the dyads are all pairs of individuals within a family (k-group round-robin design) and we wish to distinguish between the father (F), mother (M), and child (C) roles, the covariates $\boldsymbol{x}_{\alpha,a}$ could be dummy variables for the actor being the father and the actor being the mother, $\boldsymbol{x}_{\beta,p}$ dummy variables for the partner being the father and the partner being the mother (with child being the reference group). To allow the parent-as-actor and parent-as-partner effects to be different when parents interact with each other than when they interact with a child, $\boldsymbol{x}_{\gamma,a,p}$ could include the interactions between the dummy variables for the actor being the father and for the partner being the mother and between the dummy variables for the actor being the mother and for the partner being the father. This model would produce seven unconstrained means for the seven types of dyads, where (actor, partner) are (C, C), (M, C), (F, C), (C, M), (C, F), (F, M), or (M, F), as discussed in Snijders and Kenny (1999).

Keeping in mind that the response probability for actor a when combined with partner p is a function of the composite latent variable $\theta_{a,p}$, whose mean is $\mu_{\alpha,a} + \mu_{\beta,p} + \mu_{\gamma,a,p}$, care must be taken to ensure that the regression coefficients are identified. For instance, if one of the covariates for $\mu_{\gamma,a,p}$ is the difference in some attribute, $z_a - z_p$, between the actor and partner, it is not possible to also include both the attribute for the actor, z_a, in the model for $\mu_{\alpha,a}$ and the attribute for the partner, z_p, in the model for $\mu_{\beta,p}$. Another example where identification is impossible is where dyads are males paired only with females (i.e., if the actor is a male, then the partner must be a female and vice versa) and gender is included as a covariate in the models for both $\mu_{\alpha,a}$ and $\mu_{\beta,p}$. Such an example is described in greater detail in Sect. 4.

It is also possible to allow the variances of the latent traits to depend on covariates. For instance, Snijders and Kenny (1999) allow the actor and partner variances to be different for fathers, mothers, and children. In their application which included only child–parent dyads, one model also allowed the variance of $\gamma_{a,p}$ to depend on whether the parent in the dyad was the mother or the father. Such an approach allows modeling non-exchangeable dyads in general.

2.3.2. Including Random Effects for the Latent Traits If the individuals are clustered in different ways, e.g., sibling groups, schools, and/or neighborhoods, it may make sense to include cluster-level random effects to allow the actor, partner, and dyadic effects to be higher or lower, on average, in some clusters than others, or, in other words, to have intraclass correlations. An obvious specification would be to introduce corresponding cluster-level actor, partner, and dyadic effects, A_j, B_j, and G_j, respectively, for cluster j. The corresponding expression for $\theta_{a,p}$ then becomes

$$\theta_{a,p} \equiv \alpha_a + \beta_p + \gamma_{a,p} + A_{j[a]} + B_{j[p]} + \mathrm{I}(j[a] = j[p])G_{j[a]},$$

where $j[a]$ is the cluster that individual a belongs to and $\mathrm{I}(j[a] = j[p])$ is an indicator for actor and partner belonging to the same cluster. We could specify a bivariate normal distribution for A_j, B_j with variances σ_A^2, σ_B^2 and correlation ρ_{AB} and a normal distribution for G_j with variance σ_G^2.

As shown in "Appendix A," these four additional parameters are identified if individuals in the same dyad can belong to either the same cluster or to different clusters. An example of such a design would be pairing individuals within and between sibling groups (or twin pairs) to study how within-family interactions vary between families and differ from between-family interactions. Another example would be dyads formed in a professional development workshop where groups of individuals come from different institutions, for instance groups of teachers from different schools, and pairs can be formed both within and between institutions. A third example is forming the clusters/groups as part of a social experiment to see how some designed group activity affects subsequent dyadic interactions among same-group dyads compared with different-group dyads. In such designs, σ_A^2 would be interpreted as the between-cluster variance or the within-cluster covariance of the total actor effects, $\alpha_a + A_{j[a]}$. Similarly, σ_B^2 is the between-cluster variance or within-cluster covariance of the total partner effects, $\beta_p + B_{j[p]}$. The term $\sigma_A \sigma_B \rho_{AB}$ is added to the covariance between total actor and total partner effects if the actor and partner belong to the same cluster (even if they are not the same individual). Clusters where actor effects are larger than average may also have larger than average partner effects which would correspond to a positive ρ_{AB}. Finally, σ_G^2 is the additional dyadic reciprocity covariance for same-cluster dyads compared with different-cluster dyads.

As mentioned above and shown in "Appendix A," the model is identified if it is possible for individuals to be in the same dyad but not the same cluster. In contrast, if dyads are formed only among individuals within the same cluster, e.g., students are paired only with other students from the same class, "Appendix A" shows that it is not possible to estimate the individual variance and correlation parameters, and it hence makes sense to define $u_j \equiv A_j + B_j + G_j$ with one variance parameter. Such models have been used for k-group round-robin designs where the groups are naturally occurring, such as families (e.g., Snijders and Kenny 1999).

It is of course possible to handle multiple nested or non-nested classifications by adding the corresponding random intercepts u if dyads are formed only within a classification and A, B, and G if dyads are formed both within and across classifications (e.g., neighborhood when dyads are formed within schools or firms). Non-exchangeability can be handled by specifying different (co)variances for u or for A, B, and G for different groups of individuals.

2.3.3. Distal Outcomes The dIRT model can be extended by using, for instance, generalized linear models to model one or more distal outcomes, where α_a, β_p, and $\gamma_{a,p}$ are latent covariates. For example, we can consider a binary distal outcome $d_{a,p}$ of a dyad (a, p) that takes the value of 1 with the conditional probability $\pi_{a,p}$ given the latent traits, and 0 otherwise. For the speed-dating application considered in Sect. 4, the distal outcome is whether each actor in a dyad elected to see the partner again. Here, $\pi_{a,p}$ can be modeled using a logistic regression model, such as

$$\log\left(\frac{\pi_{a,p}}{1-\pi_{a,p}}\right) = b_0 + b_1\alpha_a + b_2\alpha_p + b_3\beta_a + b_4\beta_p + b_5\gamma_{a,p} + b_6\gamma_{p,a} + b_7\alpha_a\alpha_p + b_8\beta_a\beta_p + b_9\gamma_{a,p}\gamma_{p,a}. \tag{6}$$

Distal outcome regressions can also include observed covariates of both the individual and the dyad if necessary.

Notice that in the above example, the distal outcome is directed in the sense that it depends on which individual in the dyad plays the role of the actor, and the individuals are therefore not exchangeable in the sense that the effect of α_a on the distal outcome for actor a is not necessarily the same as the effect of α_p. If the distal outcome is undirected, however, and individuals within dyads are exchangeable (e.g., in the case of pairs of individuals participating in a collaborative problem-solving task where the outcome of interest is how well the task was completed per pair), then (6) should be constrained to have $b_1 = b_2$, $b_3 = b_4$, and $b_5 = b_6$. If there is one undirected outcome per dyad and the individuals in the dyad are non-exchangeable (e.g., males paired with females), such a constraint is not needed if, for instance, a represents the male and p the female in the dyad.

2.4. *Variants of the SRM*

SRMs can be used when individuals appear in multiple dyads and actor/partner role reversals (or reciprocals) occur as discussed in Sect. 2.2. Here, we give an overview of different variants of the SRM model that have been discussed in the literature.

In a seminal paper, Warner et al. (1979) proposed a round-robin model that was subsequently called the *complete* SRM by Kenny and La Voie (1984). The model has also been discussed, for instance, by Gill and Swartz (2001) and Li and Loken (2002). Wong (1982) mentions special cases. In the complete SRM, a continuous observed outcome y of actor a in the presence of partner p measured over multiple time points t is modeled as

$$y_{a,p,t} \;=\; \mu + \alpha_a + \beta_p + \gamma_{a,p} + \epsilon_{a,p,t}.$$

The distributions of the random effects α_a, β_a, $\gamma_{a,p}$ and $\gamma_{p,a}$ are identical to the distributions of the corresponding latent traits in the dIRT; see (3) and (4). The normally distributed error terms $\epsilon_{a,p,t}$ in the complete SRM are uncorrelated with the random effects and themselves uncorrelated for different t and can therefore be interpreted as test–retest measurement errors. Because $\epsilon_{a,p,t}$ reduces to $\epsilon_{a,p}$ if there is only one time point, the identifiability of the above model, and in particular of the variance of $\gamma_{a,p}$ separately from that of $\epsilon_{a,p,i}$, hinges crucially on measurements of the same dyad across multiple time points. In the dIRT model, multiple items essentially play the role of multiple time points, allowing for identification of the variance of $\gamma_{a,p}$ because of local independence.

If multiple time points or items are not available, a model denoted the *basic* SRM by Kenny et al. (2006) can be used; see also, for instance, Kenny and Kashy (1994), Lüdtke et al. (2013), Snijders and Kenny (1999), and Li and Loken (2002). The basic SRM can be obtained by assuming that the dyadic effect is symmetric across the actor/partner role individuals play, i.e., $\gamma_{a,p} = \gamma_{p,a}$ for two individuals a and p. Now $\gamma_{a,p}$ simply induces additional dependence between the two directed responses for a given (undirected) dyad. We can equivalently replace $\gamma_{a,p} + \epsilon_{a,p}$ by a single error term, typically denoted $\gamma_{a,p}$, that is correlated across members of the same dyad. The basic SRM can then be written as

$$y_{a,p} \;=\; \mu + \alpha_a + \beta_p + \gamma_{a,p}.$$

Li and Loken (2002) show the correspondence between this model and a diallel model used in genetics (e.g., Cockerham and Weir 1977).

In a k-group round-robin design, it may be useful to include a group-level random intercept, for instance, when the groups are families, with each pair of family members forming a dyad (e.g., Snijders and Kenny 1999; Loncke et al. 2018). Such a model is described at the end of Sect. 2.3.2. Also, all the models described in this section can be extended to allow means and variances to differ between groups of individuals (e.g., fathers, mothers, and children) as described in Sect. 2.3.1.

Multivariate extensions of the SRM have been proposed for the situation where actors provide ratings on several continuous variables (e.g., Nestler 2018; Lüdtke et al. 2018). For the case with a single time point, the random part of the model can be written as

$$y_{a,p,i} \;=\; \alpha_{a,i} + \beta_{p,i} + \gamma_{a,p,i}$$

for variable i, where $\epsilon_{a,p,i}$ has been removed because only one error term (correlated across members of the same dyad) can be included as in the basic SRM. Unstructured covariance matrices are specified for each of these terms across variables, and in addition to the same-variable covariances between $\alpha_{a,i}$ and $\beta_{a,i}$ and between $\gamma_{a,p,i}$ and $\gamma_{p,a,i}$ that are part of a basic univariate SRM, the model allows for all corresponding cross-variable covariances as well.

The SRM with stable *construct effects* (Kenny 1994, "Appendix B") is a more structured multivariate model that can be written as

$$y_{a,p,i} \;=\; \mu + \mu_i + \alpha_a + \alpha_{a,i} + \beta_p + \beta_{p,i} + \gamma_{a,p} + \gamma_{a,p,i},$$

where α_a, β_p, and $\gamma_{a,p}$ are stable actor, partner, and dyadic constructs, respectively that are constant across different variables or items. In contrast, $\alpha_{a,i}$, $\beta_{p,i}$, and $\gamma_{a,p,i}$ are the corresponding unstable constructs that vary over items (and $\gamma_{a,p,i}$ includes measurement error). The sum of the stable constructs corresponds to the composite latent variable in the dIRT. The SRM with stable construct effects can be viewed as including a classical test theory measurement model for the sum of all constructs, with measurement error absorbed in $\gamma_{a,p,i}$. Bonito and Kenny (2010) accommodated groups in the model by including group-specific fixed and random effects.

As far as we know, the common factor version of our dIRT model (where the measurement model for $\theta_{a,p}$ is a univariate factor model for continuous responses) has not been discussed in the literature.

Hoff (2005) [see also Hoff (2015)] proposed the additive and multiplicative effects (AME) model for dyadic data. Considering continuous responses and ignoring observed covariates, this model can be written as

$$y_{a,p} \;=\; \mu + \alpha_a + \beta_p + \mathbf{U}_a'\mathbf{V}_p + \epsilon_{a,p},$$

where $\mathbf{U}_a'\mathbf{V}_p$ is a bilinear effect of the latent component vectors $\mathbf{U}_a$ and $\mathbf{V}_p$. The AME model can be viewed as combining features of the factor analysis of variance (FANOVA) model of Gollob (1968) with the basic SRM model.

We are aware of only a few papers that extend the SRM model to handle non-continuous responses. For the basic SRM, the AME model of Hoff (2005) encompasses generalized linear models, whereas Koster and Leckie (2014) used bivariate Poisson models for counts and Koster and Brandy (2018) used bivariate probit models for binary responses.

For longitudinal data, Hoff (2015) mentioned a dynamic version of the SRM that, ignoring observed covariates, can be expressed as

$$y_{a,p,t} = \mu + b_1 y_{a,p,t-1} + b_2 y_{p,a,t-1} + \alpha_a + \beta_p + \epsilon_{a,p,t},$$

where $y_{a,p,t-1}$ and $y_{p,a,t-1}$ are lagged values of the responses. For longitudinal dyadic data with dyads nested in groups g, Nestler et al. (2017) proposed a social relations linear growth model that can be written as

$$y_{a,p,t,g} \;=\; b_{0,a,p,g} + b_{1,a,p,g} T_{a,p,t,g} + \epsilon_{a,p,t,g},$$

where $T_{a,p,t,g}$ is the time associated with occasion t for the dyad (a, p) in group g. The intercept and slope are structured as

$$b_{0,a,p,g} = \mu_0 + \zeta_{0,g} + \alpha_{0,a,g} + \beta_{0,p,g} + \delta_{0,a,p,g}$$
$$b_{1,a,p,g} = \mu_1 + \zeta_{1,g} + \alpha_{1,a,g} + \beta_{1,p,g} + \delta_{1,a,p,g}.$$

The random group effects $\zeta_{0,g}$ and $\zeta_{1,g}$ for the intercept and slope, respectively, have an unstructured covariance matrix. These effects are independent from the random actor effects $\alpha_{0,a,g}$ and $\alpha_{1,a,g}$ for the intercept and slope and the corresponding random partner effects $\beta_{0,p,g}$ and $\beta_{1,p,g}$. The random actor and partner effects have an unstructured covariance matrix.

3. Estimation

The dIRT model includes crossed random effects so that the marginal likelihood involves high-dimensional integrals. For example, in a k-group block design, the dimensionality of integration for the likelihood contribution of a group is the smaller of the two block sizes within the group plus one (Goldstein 1987). Numerical integration or Monte Carlo integration quickly becomes prohibitive, and approximate methods are often not satisfactory (see, e.g., Jeon et al. 2017 and references therein). Fortunately, Bayesian estimation via Markov-chain Monte Carlo (MCMC) is feasible, and we adopt this approach here. MCMC has previously been used by Lüdtke et al. (2013) for the basic SRM, Gill and Swartz (2001) for the complete SRM, and Hoff (2005) for the AME model. For dIRT, we use the "No-U-Turn" sampler (Hoffman and Gelman 2014) implemented in `Stan` (Stan Development Team 2018a). The `Stan` language affords us great flexibility in extending the basic dIRT model. We also verified all results using `Matlab` (version r2016b) via custom-written code based on the Metropolis–Hastings algorithm (Metropolis and Ulam 1949).

To use MCMC, we define prior distributions for the parameters in (3) and (4) as well as the item parameters in (2) and potentially the coefficients of the distal outcome regression in (6). When no prior information is available, as in the speed-dating application, diffuse priors can be used to obtain estimates that are similar to maximum likelihood estimates. For example, in the next section, we take the priors of all variances, σ^2_α, σ^2_β, and σ^2_γ, to be uniform on the interval $[0, +\infty)$ and the priors for the correlations, $\rho_{\alpha\beta}$ and ρ_γ, to be uniform on the interval $[-1, 1]$. For step difficulties $\delta_{i,j}$ and regression coefficients $b_0, b_1, \ldots, b_9$ in the distal outcome model (6), we specify uniform priors $(-\infty, +\infty)$. In Sect. 5, we also consider other diffuse priors for these parameters.

Four MCMC chains with random starting values are run for 2000 iterations, with a burn-in period of 1000 iterations. The posterior expectations of the parameters are then estimated by the sample means of the converged (post burn-in) MCMC draws for the four chains, i.e., they are based on an MCMC sample size of 4,000. Convergence is assessed by monitoring the $\hat{R}$ statistic (Gelman and Rubin 1992).

The distal outcome model in (6) can be estimated jointly with the dIRT model by combining the log-likelihood contributions from the dIRT and distal outcome models in forming the joint log posterior of all parameters, given the dIRT item responses and distal outcome. Bayesian packages typically only require explicit definition of the joint likelihood of all responses given the latent variables, model parameters, and covariates, and this joint likelihood is simply the product of the likelihoods of the dIRT and distal outcome models because of conditional independence.

Joint estimation of the dIRT and distal outcome models is consistent and asymptotically efficient if both models are correctly specified. However, to protect against misspecification of the distal outcome (or "structural") model, a sequential approach could be used where the parameters of the dIRT ("measurement") model are estimated in step 1 and subsequent steps are used to obtain estimates of the structural (distal outcome) model parameters. If the measurement model is correctly specified, the estimates from step 1 are consistent even if the structural model is misspecified. However, if the structural model is correctly specified, joint estimation is more efficient than sequential approaches. From a conceptual point of view, it has been argued in the structural equation modeling and latent class literature that altering the structural model by, for instance, adding or removing distal outcomes affects the interpretation of the latent variables because these distal outcomes play a similar role to the items or indicators that define the latent traits. Sequential modeling can protect against such "interpretational confounding" (Burt 1976) where the meaning of a construct is different from the meaning intended by the researcher [see Bakk and Kuha (2018) for further discussion].

The most obvious sequential approach is to use factor score regression (Skrondal and Laake 2001) where one estimates the measurement model (step 1), obtains judiciously chosen scores for the latent traits from the measurement model (step 2), and substitutes these scores for the latent traits to estimate the structural model as if the latent traits were observed (step 3). This approach was adopted by Loncke et al. (2018) for SRMs. However, factor score regression is only consistent for link functions that are rarely of relevance in IRT (such as the identity) and naive standard errors from this approach are moreover underestimated. To address these limitations, a multiple imputation approach can be used, where multiple draws of the latent traits are obtained from their posterior distribution and the estimates for the structural model are combined using Rubin's formula (Rubin 1987). Lüdtke et al. (2018) use such an approach in an SRM to estimate covariate effects on individual-level latent traits. Multiple imputation is natural in a Bayesian setting where full posteriors of the latent traits are available. A more straightforward pseudo-likelihood estimator, in the sense of Gong and Samaniego (1981), was proposed by Skrondal and Kuha (2012) (see also Bakk and Kuha 2018). In this case, the measurement model is first estimated, followed by joint estimation of the measurement and structural models under the constraint that the parameters of the measurement model are set equal to the estimates from the first stage. We present the results of the joint approach in this paper and include results for the sequential approach with multiple imputation as Supplementary Material.

4. Speed-Dating Application

We use a speed-dating dataset (Fisman et al. 2006) to examine the mutual attractiveness ratings of both individuals in a dyad to look for evidence of interactions that cannot be explained solely by the individuals' attractiveness or rating preferences. We also consider whether males and females differ in how they perceive their interactions. Additionally, by treating the final dating decision of whether the actor wants to see the partner again as a distal outcome, we investigate to what extent the dating decision relates to the dyadic latent trait.

The data were collected at 21 separate researcher-organized speed-dating sessions, over a period of 2 years, with 10–44 students from graduate and professional schools at Columbia University in each session. During these sessions, attended by nearly an equal number of male and female participants, all members of one gender would meet and interact with every member of the opposite gender for 5 min each. At the end of the 5-min session, participants would rate their partner based on five attractiveness items on a form attached to a clipboard that they were provided with. Each item was rated on a 10-point Likert-scale, which we collapsed to a 5-point scale by combining pairs of adjacent response categories to mitigate sparseness. Participants rated each other on five different items, all related to the overall attractiveness (viz. physical attractiveness, ambition, how fun they were, intelligence, and sincerity) of the partner, as detailed partner, p

actor, a		F_1	F_2	F_3	F_4	F_5	M_6	M_7	M_8	M_9	M_{10}
	F_1						$y_{1,6}$	$y_{1,7}$	$y_{1,8}$	$y_{1,9}$	$y_{1,10}$
	F_2						$y_{2,6}$	$y_{2,7}$	$y_{2,8}$	$y_{2,9}$	$y_{2,10}$
	F_3						$y_{3,6}$	$y_{3,7}$	$y_{3,8}$	$y_{3,9}$	$y_{3,10}$
	F_4						$y_{4,6}$	$y_{4,7}$	$y_{4,8}$	$y_{4,9}$	$y_{4,10}$
	F_5						$y_{5,6}$	$y_{5,7}$	$y_{5,8}$	$y_{5,9}$	$y_{5,10}$
	M_6	$y_{6,1}$	$y_{6,2}$	$y_{6,3}$	$y_{6,4}$	$y_{6,5}$					
	M_7	$y_{7,1}$	$y_{7,2}$	$y_{7,3}$	$y_{7,4}$	$y_{7,5}$					
	M_8	$y_{8,1}$	$y_{8,2}$	$y_{8,3}$	$y_{8,4}$	$y_{8,5}$					
	M_9	$y_{9,1}$	$y_{9,2}$	$y_{9,3}$	$y_{9,4}$	$y_{9,5}$					
	M_{10}	$y_{10,1}$	$y_{10,2}$	$y_{10,3}$	$y_{10,4}$	$y_{10,5}$					

FIGURE 2.
Example of responses $y_{a,p}$ of actor a rating partner p in a single-group block-dyadic structure consisting of five females $F_1, \ldots, F_5$ and five males $M_6, \ldots, M_{10}$.

in Supplementary Material. They also indicated whether or not they would like to see the partner again. We dropped all invalid ratings (responses that were either missing or outside the instructed range) from an actor of a partner and the corresponding ratings from the partner of the actor even if the latter was valid. This amounted to a loss of less than 5% of the data.

After data cleaning, we had a total of 551 individuals, interacting in 4184 distinct pairs, leading to 8368 surveys completed (twice the number of pairs, given that both members of a pair rated each other). This corresponds to the "k-group block-dyadic" design described in Sect. 2.2. An illustrative example of data collected for one item in a balanced group of ten individuals is shown in Fig. 2.

In the data, the rating by actor a of partner p on item i is given by $y_{a,p,i}$. In addition to each individual's rating of his/her partner, we also had access to an indicator $d_{a,p}$ for whether actor a elected to see partner p again. Note that this indicator is directional and $d_{a,p}$ may therefore differ from $d_{p,a}$. However, embedding the dIRT model within a distal outcome regression where the distal outcome is undirectional is also possible. For example, if we knew whether the dyad did in fact go on a date, this outcome would be unique to the dyad.

The joint MCMC estimation approach described in Sect. 3 was used to estimate the dIRT (2) together with the distal regression in (6) while constraining $b_7 = b_8 = b_9 = 0$. Estimates of the standard deviations and correlations of individual and dyadic latent traits are presented in Table 1, together with the estimated Monte Carlo error (MC Err), posterior standard deviation (SD), 95% credible interval (95% CI), effective sample size (Eff N), and $\hat{R}$. We also estimated two models similar to the model above: one without constraining $b_7 = b_8 = b_9 = 0$ (presented as Supplementary Material), and another with a parameter to account for the difference between the expected attractiveness of males and that of females (see Sect. 4.4). Both models produce similar estimates for the parameters in Table 1.

TABLE 1.
Estimates of standard deviations and correlations of individual and dyadic latent traits (joint approach).

	Mean	MC Err	SD	95% CI	Eff N	$\hat{R}$
σ_α	1.03	0.0010	0.0373	(0.96, 1.10)	1501	1.0013
σ_β	0.63	0.0012	0.0261	(0.58, 0.68)	450	1.0035
σ_γ	0.98	0.0009	0.0177	(0.95, 1.01)	405	1.0177
$\rho_{\alpha\beta}$	− 0.06	0.0016	0.0529	(− 0.16, 0.04)	1043	1.0007
ρ_γ	0.46	0.0011	0.0244	(0.41, 0.51)	531	1.0096

The code used to obtain the results is provided as Supplementary Material and explained in a `Stan` case study (Sim et al. 2019).

4.1. Partitioning of Variance Between Individual and Dyadic Latent Traits

Standard deviation and correlation estimates are reported in Table 1. In the dIRT, the variance of the composite latent variable $\theta_{a,p}$ is the sum of the variances of the individual- and dyad-level latent traits, α_a, β_p, and $\gamma_{a,p}$. It is instructive to examine the relative contributions of these latent traits to the composite. The percentage of the variance of $\theta_{a,p}$ that is due to α_a, β_p, and $\gamma_{a,p}$ is estimated as 44%, 16%, and 40%, respectively.

Interestingly, the variance of α_a is larger than that of β_p, implying that the actor's perception of the partner is more influenced by the actor's average tendency to rate others as attractive, which we could call actor leniency, than by the partner's average tendency to be rated as attractive, which we could think of as the partner's "universal" attractiveness. These variances are analogous to the SRM's "actor" and "partner" variances defined by Kenny et al. (2006), respectively.

While the majority (60%) of the variance is accounted for by the individual effects (α_a and β_p), the dyadic effect ($\gamma_{a,p}$) accounts for a substantial proportion of the total variance, at 40%. The analogous SRM variance is referred to as the "relationship variance" by Kenny et al. (2006). A traditional IRT model, measuring individual latent traits only, would ignore this contribution, which can be thought of as the "eye-of-the-beholder" effect. In particular, this dyadic component would not be identifiable for standard IRT data where the individual only belongs to a single dyad.

4.2. Correlations

The within-person correlation $\rho_{\alpha\beta}$ between α_a and β_a reflects the relationship between an individual's tendency to rate others as attractive and his/her own tendency to be rated as attractive. If this correlation is positive, it indicates that individuals who judge others as more attractive on average are also judged by others as more attractive on average. If negative, it indicates that individuals who judge others as less attractive on average are judged by others as more attractive on average. The corresponding correlation in the SRM is referred to as "generalized reciprocity" by Kenny et al. (2006)

The between-person correlation ρ_γ of a dyad reflects the extent to which the (directed) dyadic trait is correlated between members of a given pair. If positive, it indicates that when an individual judges their partner as more attractive, that partner also tends to judge the individual as more attractive. If negative, it suggests that members of a pair perceive each other's attractiveness in opposing ways. Kenny et al. (2006) refer to the analogous "actor–partner correlation" in the SRM as a measure of "dyadic reciprocity."

Table 1 shows that the estimate of the correlation ρ_γ is positive with a 95% credible interval that does not contain zero. In contrast, the estimate of the correlation $\rho_{\alpha\beta}$ is negative with a 95%

TABLE 2.
Estimates for Distal Outcome Regression (Joint Approach).

	Mean	MC Err	SD	95% CI	Eff N	$\hat{R}$
$b_0[1]$	− 0.88	0.0036	0.0790	(− 1.04, − 0.73)	494	1.0099
$b_1[\alpha_a]$	0.14	0.0045	0.0920	(− 0.05, 0.32)	420	1.0006
$b_2[\alpha_p]$	− 0.02	0.0014	0.0566	(− 0.13, 0.09)	1741	1.0120
$b_3[\beta_a]$	− 2.92	0.0173	0.2498	(− 3.46, − 2.49)	209	1.0134
$b_4[\beta_p]$	3.48	0.0117	0.2071	(3.12, 3.93)	312	1.0205
$b_5[\gamma_{a,p}]$	3.42	0.0161	0.2439	(3.00, 3.95)	228	1.0306
$b_6[\gamma_{p,a}]$	0.13	0.0032	0.0834	(− 0.04, 0.29)	688	1.0128

credible interval containing zero. The relatively larger estimated between-individual correlation indicates that members of each pair were likely to perceive each other's attractiveness similarly.

4.3. *Distal Outcome Regression*

Estimates of the regression coefficients of the distal outcome model (6) with the constraints $b_7 = b_8 = b_9 = 0$ are given in Table 2. (See Supplementary Material for estimates of the unconstrained model.) We see that the estimated distal outcome regression coefficients are largest, in absolute value, for: a) the individual attractiveness α of both the actor ($\hat{b}_3$) and the partner ($\hat{b}_4$), and b) the unique relationship of the dyad γ from the actor's perspective ($\hat{b}_5$) but *not* for that from the partner's perspective ($\hat{b}_6$). Finding b) is consistent with our expectations given that the distal outcome reflects the viewpoint of the actor, rather than that of the partner. However, a less obvious finding is a) because the rater's own attractiveness, $\hat{b}_3$, *negatively* influences their dating decision. This suggests that the more attractive a rater was, the less likely they were to want to see the partner again. The estimated coefficients are tiny for the leniency of both the actor and the partner, as well as for the unique relationship of the dyad from the perspective of the partner, and have 95% credible intervals either containing zero or having one limit close to zero.

4.4. *Gender Differences*

Both the basic dIRT model and the model with a distal outcome can be extended to account for differences in the way females and males perceived their social interactions. In (3), we assumed that male and female participants shared the same expected leniency μ_α and attractiveness μ_β, by setting both of these expectations to zero. We can relax this by allowing the genders to have a different expectation for one of these parameters whilst setting the other to zero. The distribution for α and β becomes:

$$\begin{bmatrix} \alpha_a \\ \beta_a \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} m_a \mu_{\text{male}} \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_\alpha^2 & \rho_{\alpha\beta}\sigma_\alpha\sigma_\beta \\ \rho_{\alpha\beta}\sigma_\alpha\sigma_\beta & \sigma_\beta^2 \end{bmatrix}\right). \tag{7}$$

Here, μ_{male} is the difference between the expected attractiveness of males and that of females, and m_a is an indicator for whether individual a is male. This gender parameter can also be interpreted as the difference between the expected leniency of females and that of males. Hence, a positive μ_{male} would suggest that males were on average more attractive than females, and/or females were more lenient in their ratings of males. We note that these effects could be disentangled if males rated other males and females rated other females. However, because we do not have such data,

μ_{male} can only be interpreted as a linear combination (with unknown constants) of the average additional male attractiveness and average additional female rater leniency.

The gender difference μ_{male} is estimated to be 0.08 with a 95% credible interval containing zero, while the other parameter estimates change negligibly. There is therefore insufficient evidence to suggest a gender difference. The variance and correlation estimates are virtually the same for the models with and without μ_{male}.

5. Simulations

We first present the results of a simulation study exploring Bayesian properties of the MCMC estimator for an extended dIRT model that includes a distal outcome. We generated data for the same data design, size, and parameter estimates as in the previous section. Starting with the estimated values of the variance and correlation hyperparameters, we generated 551 pairs of individual latent traits (α_a, β_a), and 8126 directed dyadic latent traits $\gamma_{a,p}$. Using the estimated item step difficulties from Sect. 4, we then generated responses from the dIRT model (2). Using the estimated regression coefficients, we finally generated the distal outcomes according to model (6). We summarize our findings regarding parameter recovery in the figures below.

Figures 3, 4, and 5 depict the difference between the estimated and actual parameters across all 4000 draws after convergence. In estimating the model, we considered two different sets of diffuse priors: (1) the same priors as in Sect. 4 and (2) Cauchy priors with location parameter 0 and scale parameter 5 for all standard deviations (suggested in Gelman 2006), LKJ priors with shape parameter 1.1 for the correlation matrices of the individual and dyadic traits (recommended in Stan Development Team 2018b), and normal priors with mean 0 and standard deviation 7 for the step difficulties and regression coefficients. The squares represent the posterior means, while the whiskers represent the bounds for the 95% credible intervals based on the 2.5th and 97.5th percentiles. We see that the credible intervals for all but one of the 32 parameters contain the true value, and our procedure hence has good Bayesian performance. Furthermore, the different choices of diffuse priors do not affect the eventual estimates and credible intervals much.

In order to evaluate frequentist properties such as the bias of point estimates and the validity of model-based standard errors, we generated 50 datasets based on the same procedure as above and estimated the same model for each dataset. Based on these 50 replications, we then estimated (i) the absolute bias of parameter estimates using the difference between the mean (over replications) of the estimated parameters and the true values, and (ii) the relative error of standard error estimates using the mean (over replications) of the estimated standard errors divided by the empirical standard deviation (over replications) of the point estimates minus 1. Monte Carlo errors for these quantities were estimated using the formulae in White (2010).

In Fig. 6, we show the estimated absolute bias of the parameter estimates (top) and relative error of the standard error estimates (bottom), together with error bars of ± 1.96 times their Monte Carlo error estimates. The intervals represent approximate 95% confidence intervals if the sampling distributions are approximately normal. We see that there is small absolute bias in our point estimates across parameters, most of which can be attributed to Monte Carlo error with the exception of b_2, b_4 and b_5. There is also small relative bias for the standard error estimates, most of which can be attributed to Monte Carlo error with the exception of $\delta_{2,2}$ and $\delta_{3,4}$. In summary, our procedure has good frequentist properties.

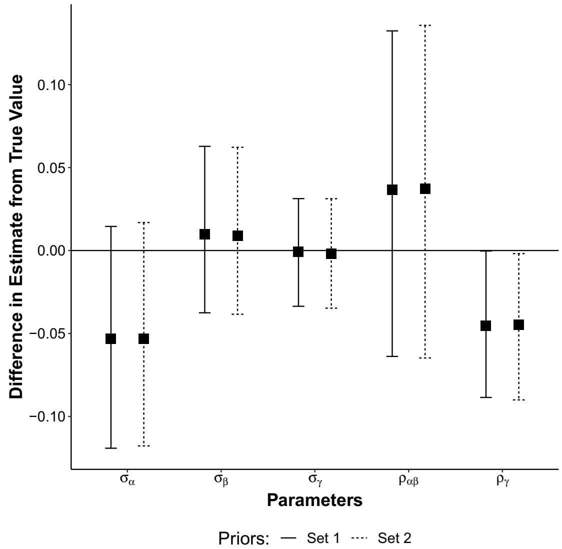

FIGURE 3.
Difference between hyperparameter estimates and true parameter values.

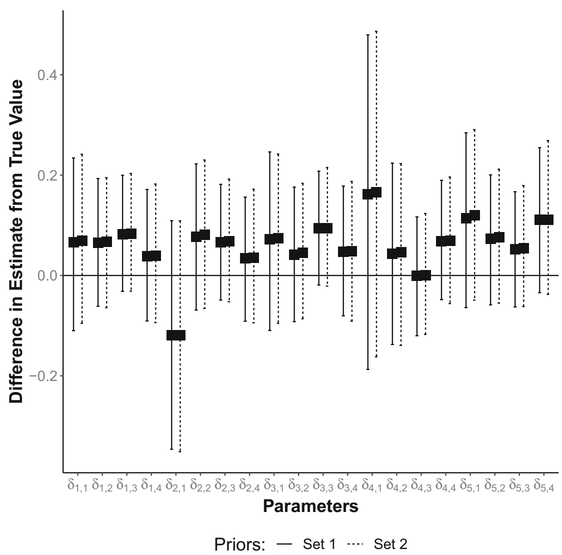

FIGURE 4.
Difference between item step difficulty estimates and true parameter values.

6. Concluding Remarks

We have proposed a dyadic Item Response Theory (dIRT) model that integrates Item Response Theory (IRT) models for measurement and the complete Social Relations Model (SRM) for dyadic data by modeling the responses of an actor as a function of the actor's inclination to act and the partner's tendency to elicit that action as well as the unique relationship of the pair. We described how the model can be extended to triads or larger group settings, include covariates for the

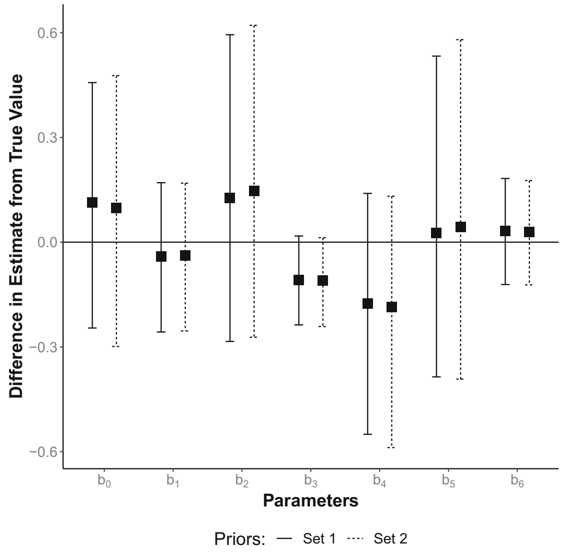

FIGURE 5.
Difference between distal outcome regression estimates and true parameter values.

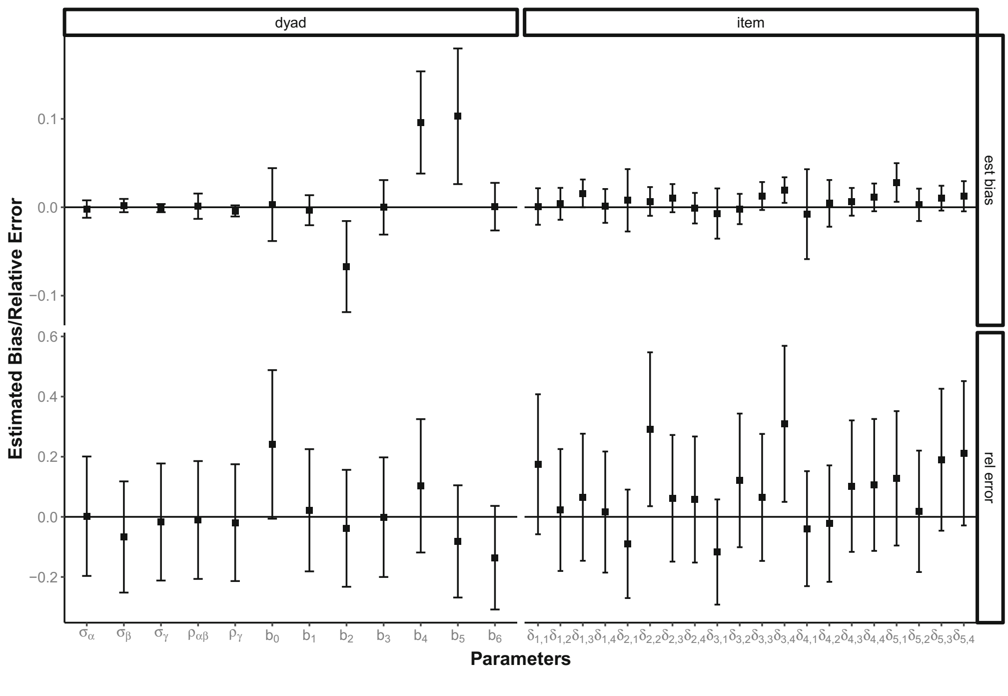

FIGURE 6.
Performance of point estimates and standard errors across 50 replications.

individual and dyad, include cluster-level random effects, and accommodate distal outcomes. We also discussed data designs for which the dIRT model is identified, emphasizing that longitudinal data are not required, and described how the model can be estimated using standard software for Bayesian inference. The proposed estimation approach was shown to have good performance in simulation studies.

The practical utility of the dIRT model was demonstrated by applying it to speed-dating data with ordinal items. The estimated variance of the actor effect suggests that there was some variation in the way different individuals rated the same sets of partners or in other words that there was variation in how lenient individuals were in rating their partners. The estimated variance of the partner effect can be thought of as reflecting how attractive the partner is, on average, to all other individuals, and indicates that there is some degree of universal attractiveness. We found that there is evidence for dyadic latent traits (unique interaction effects) and that the dyadic latent trait, from the perspective of the actor, helps predict whether the actor wants to see the partner again. This finding suggests that the dyadic latent trait has predictive validity, a conclusion that can perhaps be more easily justified when a sequential estimation approach is used. A traditional IRT model, measuring individual latent traits only, would ignore this dyadic latent trait, which can be thought of as the "eye-of-the-beholder" effect. The dyadic latent traits were positively correlated within dyads, suggesting that both members of a dyad tended to perceive their interaction similarly.

In the speed-dating application, the dyadic latent trait was of particular interest from the point of view of matchmaking. In other applications where the actors can be viewed as the raters, "perceivers," or informants used to make inferences regarding the partners, the partner latent trait is of greatest interest. In this case, the advantage of the dIRT is that it purges the measurement of the partner latent trait from both the global rater bias α and the target-specific rater bias γ. In a collaborative problem-solving task, both the actor and partner latent traits may be of interest, in which case it becomes important to accommodate the dyadic latent trait in the model to prevent it from contaminating the individual latent traits of interest. The dyadic latent trait could in this case be viewed as a nuisance reflecting a fortunate or unfortunate choice of collaborator. For all these types of applications, dyadic designs that permit estimation of the dIRT are essential.

The formulation of the dIRT model, and providing a viable estimation approach for it, provides researchers with the impetus to collect appropriate data for investigating dyadic interactions or individual latent traits, free from such interaction effects, in a measurement context.

Acknowledgments

Brian Gin and Nicholas Sim are joint first authors of this paper. This work was partly funded by the Research Council of Norway through its Centres of Excellence funding scheme, project number 262700, and by the Institute of Education Sciences, U.S. Department of Education, through Grants R305D140059 and R305D190048.

Appendix A: Identifiability of Model with Random Effects

To see that the parameters σ_A^2, σ_B^2, ρ_{AB}, and σ_G^2 of the model in Sect. 2.3.2 are identified, consider $\mathrm{cov}(\theta_{a,p},\theta_{b,q})$ when a, p, b, q are four different individuals:

$$\begin{aligned}\mathrm{cov}(\theta_{a,p},\theta_{b,q}) &= E[(A_{j[a]}+B_{j[p]}+I(j[a]=j[p])G_{j[a]})(A_{j[b]}+B_{j[q]}+I(j[b]=j[q])G_{j[b]})]\\ &= I(j[a]=j[b])\sigma_A^2+I(j[p]=j[q])\sigma_B^2+[I(j[a]=j[q])+I(j[b]=j[p])]\rho_{AB}\sigma_A\sigma_B\\ &\quad + I(j[a]=j[p]=j[b]=j[q])\sigma_G^2\end{aligned}$$

This corresponds to seven distinct expressions, depending on whose cluster membership is shared:

1. The four individuals (a, p, b, q) are in four different clusters: $\mathrm{cov}(\theta_{a,p},\theta_{b,q})=0$
2. They are in three clusters, with both actors in the same cluster
   Iff $j[a]=j[b]$, $j[a]\neq j[p]$, $j[a]\neq j[q]$, $j[p]\neq j[q]$: $\mathrm{cov}(\theta_{a,p},\theta_{b,q})=\sigma_A^2$
3. They are in three clusters, with both partners in the same cluster
   Iff $j[p]=[q]$, $j[a]\neq j[p]$, $j[b]\neq j[p]$, $j[a]\neq j[b]$: $\mathrm{cov}(\theta_{a,p},\theta_{b,q})=\sigma_B^2$
4. They are in two clusters, with both actors together and both partners together
   Iff $j[a]=j[b]$, $j[p]=j[q]$, $j[a]\neq j[p]$: $\mathrm{cov}(\theta_{a,p},\theta_{b,q})=\sigma_A^2+\sigma_B^2$
5. They are in three clusters, with actor of one dyad in same cluster as partner of other dyad
   Iff $(j[a]=j[q], j[a]\neq [p], j[a]\neq j[b], j[b]\neq j[p])$
   OR $(j[b]=j[p], j[b]\neq [q], j[b]\neq j[a], j[a]\neq j[q])$: $\mathrm{cov}(\theta_{a,p},\theta_{b,q})=\rho_{AB}\sigma_A\sigma_B$
6. They are in two clusters, each with actor of one dyad and partner of the other dyad
   Iff $j[a]=j[q]$, $j[b]=j[p]$, $j[a]\neq j[p]$: $\mathrm{cov}(\theta_{a,p},\theta_{b,q})=2\rho_{AB}\sigma_A\sigma_B$
7. They are in one cluster
   Iff $j[a]=j[p]=j[b]=j[q]$: $\mathrm{cov}(\theta_{a,p},\theta_{b,q})=\sigma_A^2+\sigma_B^2+2\rho_{AB}\sigma_A\sigma_B+\sigma_G^2$

Solving Eqs. 2, 3, 6, and 7 identifies σ_A^2, σ_B^2, ρ_{AB}, and σ_G^2. The additional equations above show that some of these parameters are overidentified. To identify the remaining parameters, σ_α^2, σ_β^2, σ_γ^2, $\rho_{\alpha\beta}$, and ρ_γ, we can add the following equations for four different individuals, a, p, b, q, that belong to four different clusters:

8. $\mathrm{var}(\theta_{a,p})=\sigma_\alpha^2+\sigma_\beta^2+\sigma_\gamma^2+\sigma_A^2+\sigma_B^2$
9. $\mathrm{cov}(\theta_{a,p},\theta_{p,a})=2\rho_{\alpha\beta}\sigma_\alpha\sigma_\beta+\rho_\gamma\sigma_\gamma^2+2\rho_{AB}\sigma_A\sigma_B$
10. $\mathrm{cov}(\theta_{a,p},\theta_{a,q})=\sigma_\alpha^2+\sigma_A^2$
11. $\mathrm{cov}(\theta_{a,p},\theta_{b,p})=\sigma_\beta^2+\sigma_B^2$
12. $\mathrm{cov}(\theta_{a,p},\theta_{b,a})=\rho_{\alpha\beta}\sigma_\alpha\sigma_\beta+\rho_{AB}\sigma_A\sigma_B$

If clusters of size four do not exist (e.g., clusters are twin pairs), Eq. 7 cannot be used to identify σ_G^2. In this case, use either of the following:

$$\mathrm{var}(\theta_{a,p})=\sigma_\alpha^2+\sigma_\beta^2+\sigma_\gamma^2+\sigma_A^2+\sigma_B^2+I(j[a]=j[p])(2\rho_{AB}\sigma_A\sigma_B+\sigma_G^2)$$

or

$$\mathrm{cov}(\theta_{a,p},\theta_{p,a})=2\rho_{\alpha\beta}\sigma_\alpha\sigma_\beta+\rho_\gamma\sigma_\gamma^2+2\rho_{AB}\sigma_A\sigma_B+I(j[a]=j[p])(\sigma_A^2+\sigma_B^2+\sigma_G^2).$$

If dyads can only be formed within clusters, $\sigma_A^2+\sigma_B^2+2\rho_{AB}\sigma_A\sigma_B+\sigma_G^2$ appears in all variances and covariances except for $\mathrm{cov}(\theta_{a,p},\theta_{b,q})$ if dyad (a, p) is not in the same cluster as dyad (b, q). This can occur only if the two dyads do not share any individuals in common, in which case we obtain $\text{cov}(\theta_{a,p}, \theta_{b,q}) = \sigma_A^2 + \sigma_B^2 + 2\rho_{AB}\sigma_A\sigma_B + \sigma_G^2$ if dyad (a, p) belongs to the same cluster as dyad (b, q) and $\text{cov}(\theta_{a,p}, \theta_{b,q}) = 0$, otherwise. It follows that only the sum $\sigma_A^2+\sigma_B^2+2\rho_{AB}\sigma_A\sigma_B+\sigma_G^2$ is identified and therefore it makes sense to define $u_j \equiv A_j+B_j+G_j$, with one variance parameter σ_u^2, and to include u_j directly in the model for $\theta_{a,p}$.

References

Ackerman, R. A., Kashy, D. A., & Corretti, C. A. (2015). A tutorial on analyzing data from speed-dating studies with heterosexual dyads. *Personal Relationships*, *22*, 92–110.

Alexandrowicz, R. W. (2015). Analyzing dyadic data with IRT models. In M. Stemmler, A. von Eye, & W. Wiedermann (Eds.), *Dependent data in social sciences research* (pp. 173–202). New York: Springer.

Back, M. D., & Kenny, D. A. (2010). The social relations model: How to understand dyadic processes. *Social and Personality Psychology Compass*, *4*, 855–870.

Bagozzi, R. P., & Ascione, F. J. (2005). Inter-role relationshis in hospital-based pharmacy and therapeutics committee decision making. *Journal of Health Psychology*, *10*, 45–64.

Bakk, Z., & Kuha, J. (2018). Two-step estimation of models between latent classes and external variables. *Psychometrika*, *83*, 871–892.

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

Burt, R. S. (1976). Interpretational confounding of unobserved variables in structural equation models. *Sociological Methods & Research*, *5*, 3–52.

Card, N. A., Little, T. D., & Selig, J. P. (2008). Using the bivariate social relations model to study dyadic relationships: Early adolescents' perceptions of friends' aggression and prosocial behavior. In N. A. Card, T. D. Little, & J. P. Selig (Eds.), *Modeling dyadic and interdependent data in the developmental and behavioral sciences* (pp. 245–276). New York: Routledge.

Christensen, P. N., & Kashy, D. A. (1998). Perceptions of and by lonely people in initial social interactions. *Personality and Social Psychology Bulletin*, *24*, 322–329.

Cockerham, C. C., & Weir, B. S. (1977). Quadratic analysis of reciprocal crosses. *Biometrics*, *33*, 187–203.

De Boeck, P., & Wilson, M. (Eds.). (2004). *Explanatory item response models: A generalized linear and nonlinear approach*. New York: Springer.

Dorff, C., & Ward, M. D. (2013). Networks, dyads, and the social relations model. *Political Science Research and Methods*, *1*, 159–178.

Fisman, R., Iyengar, S. S., Kamenica, E., & Simonson, I. (2006). Gender differences in mate selection: Evidence from a speed dating experiment. *The Quarterly Journal of Economics*, *121*, 673–697.

Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis*, *1*, 515–533.

Gelman, A., & Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, *7*, 457–472.

Gill, P. S., & Swartz, T. B. (2001). Statistical analyses of round robin interaction data. *Canadian Journal of Statistics*, *82*, 321–331.

Goldstein, H. (1987). Multilevel variance components models. *Biometrika*, *74*, 430–431.

Gollob, H. F. (1968). A statistical model which combines features of factor analytic and analysis of variance techniques. *Psychometrika*, *33*, 73–115.

Gong, G., & Samaniego, F. J. (1981). Pseudo maximum likelihood estimation: Theory and applications. *The Annals of Statistics*, *74*, 861–869.

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

Hoff, P. D. (2015). Dyadic data analysis with amen. arXiv:1506.08237.

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

Jeon, M., Rijmen, F., & Rabe-Hesketh, S. (2017). A variational maximization-maximization algorithm for generalized linear mixed models with crossed random effects. *Psychometrika*, *82*, 693–716.

Kenny, D. A. (1994). *Interpersonal perception: A social relations analysis*. New York: Guilford.

Kenny, D. A. (1996). Models of nonindependence in dyadic research. *Journal of Social and Personal Relationships*, *13*, 279–294.

Kenny, D. A., & Kashy, D. A. (1994). Enhanced co-orientation in the perception of friends: A social relations analysis. *Journal of Personality and Social Psychology*, *67*, 1024–1033.

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

Kenny, D. A., & La Voie, L. (1984). The social relations model. *Advances in Experimental Social Psychology*, *18*, 141–182.

Koster, J. M., & Brandy, A. (2018). The effects of individual status and group performance on network ties among teammates in the National Basketball Association. *PLoS ONE*, *13*, e0196013.

Koster, J. M., & Leckie, G. (2014). Food sharing networks in lowland Nicaragua: An application of the social relations model to count data. *Social Networks*, *38*, 100–110.

Li, H., & Loken, E. (2002). A unified theory of statistical analysis and inference for variance components models for dyadic data. *Statistica Sinica*, *12*, 519–535.

Loncke, J., Eichelsheim, V. I., Branje, S. J. T., Buysse, A., Meeus, W. H. J., & Loeys, T. (2018). Factor score regression with social relations model components: A case study exploring antecedents and consequences of perceived support in families. *Frontiers in Psychology*, *9*, 1699.

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*, 101–119.

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

Malloy, T. E., & Kenny, D. A. (1986). The social relations model: An integrative method for personality research. *Journal of Personality*, *54*, 199–225.

Masters, G. N. (1982). A Rasch model for partial credit scoring. *Psychometrika*, *47*, 149–174.

Metropolis, N., & Ulam, S. (1949). The Monte Carlo method. *Journal of the American Statistical Association*, *44*, 335–341.

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

Nestler, S., Geukes, K., Hutteman, R., & Back, M. D. (2017). Tackling longitudinal round-robin data: A social relations growth model. *Psychometrika*, *82*, 1162–1181.

Rubin, D. B. (1987). *Multiple imputation for nonresponse in surveys*. New York: Wiley.

Sim, N., Gin, B., Skrondal, A., & Rabe-Hesketh, S. (2019). A dyadic item response theory model: Stan case study. Stan Case Studies, 6. https://mc-stan.org/users/documentation/case-studies/dyadic_irt_model.html.

Skrondal, A., & Kuha, J. (2012). Improved regression calibration. *Psychometrika*, *77*, 649–669.

Skrondal, A., & Laake, P. (2001). Regression among factor scores. *Psychometrika*, *66*, 563–575.

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

Stan Development Team. (2018a). RStan: The R interface to Stan. *R package version*, *2*(17), 3.

Stan Development Team. (2018b). Stan modeling language users guide and reference manual, version 2.18.0.

van der Linden, W. J. (Ed.). (2016). *Handbook of item response theory. Volume one: Models*. Boca Raton: CRC Press.

Warner, R., Kenny, D. A., & Stoto, M. (1979). A new round robin analysis of variance for social interaction data. *Journal of Personality and Social Psychology*, *37*, 1742–1757.

White, I. R. (2010). simsum: Analysis of simulation studies including Monte Carlo error. *The Stata Journal*, *10*, 369–385.

Wong, G. Y. (1982). Round robin analysis of variance via maximum likelihood. *Journal of the American Statistical Association*, *77*, 714–724.

Manuscript Received: 18 JUN 2019
Published Online Date: 27 AUG 2020