Score-based tests for detecting heterogeneity in linear mixed models

Ting Wang¹ · Edgar C. Merkle² · Joaquin A. Anguera³ · Brandon M. Turner⁴

Published online: 14 July 2020

© The Psychonomic Society, Inc. 2020

Abstract

Cross-level interactions among fixed effects in linear mixed models (also known as multilevel models) can be complicated by heterogeneity stemming from random effects and residuals. When heterogeneity is present, tests of fixed effects (including cross-level interaction terms) are subject to inflated type I or type II error. While the impact of variance change/heterogeneity has been noticed in the literature, few methods have been proposed to detect this heterogeneity in a simple, systematic way. In addition, when heterogeneity among clusters is detected, researchers often wish to know which clusters' variances differed from the others. In this study, we utilize a recently proposed family of score-based tests to distinguish between cross-level interactions and heterogeneity in variance components, also providing information about specific clusters that exhibit heterogeneity. These score-based tests only require estimation of the null model (when variance homogeneity is assumed to hold), and they have been previously applied to psychometric models to detect measurement invariance. In this paper, we extend the tests to linear mixed models, allowing us to use the tests to differentiate between interaction and heterogeneity. We detail the tests' implementation and performance via simulation, provide an empirical example of the tests' use in practice, and provide code illustrating the tests' general application.

Keywords Score-based tests · Heterogeneity · Linear mixed models

Introduction

Cross-level interactions are often of interest for researchers adopting linear mixed models (also known as multilevel models or hierarchical linear models), which specifically refers to the interaction between a lower-level (level 1) covariate and an upper-level (level 2) covariate. For example, in a cross-sectional study of students within schools, we might observe a cross-level interaction between students' cognitive ability (level 1) and school climate (level 2). In a longitudinal study of repeated observations within participants, we might note a cross-level interaction between time (level 1) and participant socioeconomic status (level 2). However, the estimation and testing of cross-level interactions in linear mixed models (LMMs) is often complicated

by the multiple variance terms in the model. For example, when a cross-level interaction exists, it may not be detected due to heterogeneity in a random effect variance or in a residual variance (Halaby, 2004). If this heterogeneity is not accounted for, it can lead to biased standard error estimates (Leckie, French, Charlton, & Browne, 2014; Verbeke & Lesaffre, 1996) and incorrect significance tests for fixed parameters (Kwok et al., 2007).

Heterogeneous (co-)variances often occur in longitudinal studies, where the heterogeneity in variance/covariance is observed across individuals. For example, recent ecological momentary assessment (EMA) studies have shown that substantial heterogeneity exists across individuals in the variance and covariance of emotional states over time (Röcke, Li, & Smith, 2009; Thompson et al., 2012; Knight, 2013; Ebner-Priemer et al., 2015). A similar phenomenon occurs in education research, where it is of interest to model student achievement over time (Lockwood, McCaffrey, et al., 2007). This heterogeneity is often due to unobserved covariates or to confounding with other covariates (such as time; Bresin, 2014; Koval & Kuppens, 2012); issues that are inevitable in many applied scenarios (Lockwood et al., 2007). Hence, it is important to develop accessible methods to detect heterogeneity in the multiple variance parameters observed in mixed models.

✉ Ting Wang
twb8d@mail.missouri.edu

¹ The American Board of Anesthesiology, Raleigh, NC, USA

² University of Missouri, Columbia, MO, USA

³ University of California, San Francisco, CA, USA

⁴ The Ohio State University, Columbus, OH, USA

Several previous methods test for heterogeneity by building it into the model and using a likelihood ratio (LR) test or Kolmogorov–Smirnov test (Verbeke & Lesaffre, 1996; McLachlan & Basford, 1988; Stephens, 1974). The LR test can be used when heterogeneity can be explained by observed variables. For example, *semtree* (Brandmaier, von Oertzen, McArdle, & Lindenberger, 2013), *longRpart2* (Stegmann, Jacobucci, Serang, & Grimm, 2018), and Abdollell, LeBlanc, Stephens, & Harrison (2002) all utilize a series of LR tests to detect potential split point(s) along auxiliary variables. However, this test can be cumbersome when the variable has many categories, and it can be suboptimal when the variable is ordinal or continuous. In contrast, the Kolmogorov–Smirnov test utilizes a mixture model framework, and checks the correct number of mixture components in an omnibus goodness-of-fit test. However, as shown by Verbeke & Lesaffre (1996), this approach is subject to low power, even decreasing to zero when the residual variance is large.

In this paper, we aim to extend a family of score-based tests (e.g., Merkle & Zeileis, 2013; Zeileis & Hornik, 2007) to the study of heterogeneity in linear mixed models. These tests have been previously applied to detect measurement invariance in factor analysis (Merkle & Zeileis, 2013; Merkle, Fan, & Zeileis, 2014; Wang, Merkle, & Zeileis, 2014) and in models from item response theory (IRT; e.g., Abou El-Komboz, Zeileis, & Strobl, 2014; Wang, Strobl, Zeileis, & Merkle, 2018). This study serves as a novel application of the tests in the context of LMMs by providing a unified, new approach to differentiate heterogeneity in variance components (either random effects or residual) from cross-level interactions. From a technical perspective, this study differs from the previous measurement invariance studies in that the case-wise scores are by definition no longer independent and identically distributed (i.i.d.), since the observations within the same cluster are correlated. In addition, we focus on heterogeneity in both fixed parameters and variance/covariance parameters, whereas the related approach of Fokkema, Smits, Zeileis, Hothorn, & Kelderman (2018) has tested only for changes in fixed parameters while maintaining homogeneity in variance and covariance parameters. We also discuss graphical methods associated with the tests, which can be helpful for identifying clusters that exhibit similar variance estimates. In the following section, we provide a brief overview of the score-based tests' generalizations to linear mixed models. Next, we report on the results of a simulation to examine the tests' abilities in the context of linear mixed models. Finally, we provide an empirical example with illustrating code and discuss the tests' future generalizations.

Linear mixed model

The linear mixed model (LMM) can be expressed in both conditional and marginal forms. The former facilitates theoretical understanding, and the latter simplifies the computational expression. We will detail these two expressions in the following sections, focusing on a two-level model where individual observations are nested within a series of clusters.

Conditional expression

The conditional version of the LMM can be written as

$$\mathbf{y}_j | \mathbf{b}_j \sim N(\mathbf{X}_j \boldsymbol{\beta} + \mathbf{Z}_j \mathbf{b}_j, \mathbf{R}_j) \quad (1)$$

$$\mathbf{b}_j \sim N(\mathbf{0}, \mathbf{D}) \quad (2)$$

$$\mathbf{R}_j = \sigma_r^2 \mathbf{I}_{n_j}, \quad (3)$$

where $\mathbf{y}_j$ is the observed data vector for the j th cluster, $j = 1, \dots, J$ (so that the level 1 sample size is given as $n = \sum_{j=1}^J n_j$); $\mathbf{X}_j$ is an $n_j \times p$ matrix of fixed covariates; $\boldsymbol{\beta}$ is the fixed effect vector of length p ; $\mathbf{Z}_j$ is an $n_j \times q$ design matrix of random effects; and $\mathbf{b}_j$ is the random effect vector of length q .

The vector $\mathbf{b}_j$ is assumed to follow a normal distribution with mean $\mathbf{0}$ and covariance matrix $\mathbf{D}$, where $\mathbf{D}$ is a matrix composed of variances/covariances for random effect parameters. The residual covariance matrix, $\mathbf{R}_j$, is the product of the residual variance σ_r^2 and an identity matrix of dimension n_j . Later, the matrix $\mathbf{R}$ will include residuals across all clusters, so the identity matrix is of dimension n .

Based on the notations above, the following notation is used to represent data and parameters across all clusters in the data.

$$\mathbf{y} = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_j, \dots, \mathbf{y}_J\} \quad (4)$$

$$\mathbf{X} = \{\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_j, \dots, \mathbf{X}_J\} \quad (5)$$

$$\mathbf{Z} = \{\mathbf{Z}_1, \mathbf{Z}_2, \dots, \mathbf{Z}_j, \dots, \mathbf{Z}_J\} \quad (6)$$

$$\mathbf{b} = \{\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_j, \dots, \mathbf{b}_J\} \quad (7)$$

$$\mathbf{G} = \mathbf{D} \otimes \mathbf{I}_J, \quad (8)$$

where $\otimes$ is the Kronecker product.

Finally, we define $\boldsymbol{\sigma}^2$ to be a vector of length K , containing all variance/covariance parameters (including those of the random effects and the residual). This implies that the matrix $\mathbf{D}$ has $(K-1)$ unique elements. For example, in a model with two random effects that are allowed to covary, $\boldsymbol{\sigma}^2$ is a vector of length 4 (i.e., $K=4$). The first three elements correspond to the unique entries of $\mathbf{D}$, which are commonly expressed as σ_0^2 , σ_{01} , and σ_1^2 . The last component is then the residual variance σ_r^2 .

Marginal expression

Based on Eqs. 1, 2, and 3, the marginal distribution of the LMM is

$$y_j \sim N(X_j \beta, V_j), \quad (9)$$

where

$$V_j = Z_j D Z_j^\top + R_j, \quad (10)$$

where $\top$ denotes a matrix transpose. By using the combined notation from Eqs. 4 to 7, we can further define V as

$$V = V_j \otimes I_J \quad (11)$$

$$= ZGZ^\top + R. \quad (12)$$

Thus, Eq. 9 can be rewritten as:

$$y \sim N(X\beta, V). \quad (13)$$

From Eq. 13, we can perceive the LMM as a regular linear model with correlated residual variance V . From this perspective, one can easily deduce that heterogeneity in V has little impact on the estimate of β , because $\hat{\beta}$ is still equal to $(X^\top X)^{-1} X^\top Y$, but can have a large impact on the significance test of β (Bates, Kliegl, Vasishth, & Baayen, 2015). We will illustrate this issue in the following section.

Problems stemming from heterogeneity

We now illustrate implications of heterogeneity via both theoretical results and simulation.

Theoretical demonstration

The variance-covariance matrix w.r.t. the fixed parameter corresponds to the inverse of the model's Fisher information, the relevant part of which can be expressed as (e.g., Wang & Merkle, 2018):

$$V_\beta = (XV^{-1}X^\top)^{-1} \quad (14)$$

$$= (X(ZGZ^\top + R)^{-1}X^\top)^{-1} \quad (15)$$

$$= X^{-1}(ZGZ^\top + R)(X^\top)^{-1}. \quad (16)$$

The standard error of fixed parameters, SE_β , is then the square root of the diagonal elements of V_β . This shows that V directly contributes to the fixed parameters' standard errors, which in turn influences the fixed parameters' test statistics. With the under/overestimates of SE_β , the t -statistic will be larger/smaller than it should be. Generally, one can expect that the increasing of V results in type II error, whereas decreasing

of V leads to type I error. In practice, the former happens more often. (Kwok, West, & Green, 2007) conducted a series of Monte Carlo simulations and found underspecification and misspecification of V result in overestimation of SE_β , which leads to lower statistical power in significance tests of the fixed parameters. Although their simulations only examined main effects, one can expect similar results for interaction effects. We illustrate this issue in the next section.

Data demonstration

In this section, we specifically illustrate how the change (increase) in V could impact the significance of fixed parameters by using artificial data similar to the *sleepstudy* data (Belenky et al., 2003) included in *lme4*. This dataset includes 18 subjects participating in a sleep deprivation study, where each subject's reaction time (RT)¹ was monitored for 10 consecutive days. The reaction times are nested within subject and continuous in measurement. Then we fit a model with day of measurement ("Days") as the covariate, including random intercepts and slopes that are allowed to covary. This leads to a model whose free parameters include: the fixed intercept and slope β_0 and β_1 ; the random variance and covariances σ_0^2 , σ_1^2 , and σ_{01} ; and the residual variance σ_r^2 . To illustrate the impact of heterogeneity on cross-level interactions, we also simulate an ordinal variable with four levels loosely called *Cognitive Ability* (CA), with its own main effect coefficient as β_2 and its interaction effect *Cognitive Ability* (CA) $\times$ *Days* coefficient as β_3 . In the simulation, we focus on the significance test results of β_1 and β_3 . The true values were set to be 10.47 and 6.27, respectively, with both far different from 0. The random effect variance/covariance and residual variance were set to be the same as the estimates obtained from the *sleep study* data. This leads to the model displayed from Eqs. 17 to 19.

$$RT_j | \text{Subject } j \sim N(\beta_0 + \beta_1 \text{Days} + \beta_2 \text{CA} + \beta_3 \text{Days} \times \text{CA}, R_j) \quad (17)$$

$$\text{Subject } j \sim N\left(0, \begin{bmatrix} \sigma_0^2 & \sigma_{01} \\ \sigma_{01} & \sigma_1^2 \end{bmatrix}\right) \quad (18)$$

$$R_j = \sigma_r^2 I_{10} \quad (19)$$

From Eq. 10, it can be observed that changes in V can come from either G , which is composed of between-subjects variance parameters σ_0 , σ_1 and σ_{01} , or the residual variance σ_r^2 . We generated data so that V changed with each of these four parameters, including the between-subjects intercept variance σ_0^2 , slope variance σ_1^2 , covariance σ_{01} , and residual variance σ_r^2 , along with different sample sizes as small ($n = 120$), medium ($n = 480$), and large ($n = 960$). Changes in these

¹ Strictly speaking, the response variable should be $\log(\text{RT})$ to have a meaningful infinity support.

variance parameters began at cognitive ability level 2 and were consistent thereafter. Participants below cognitive ability level 2 deviated from participants at or above level 2 by d times the parameters' asymptotic standard errors (scaled by $\sqrt{n}$), with $d=0, 1, 2, 3, 4$. To obtain the asymptotic standard error, we first fit a model under the above parameter settings but with a large sample size, e.g., 9600; then the asymptotic standard error can be extracted by taking the square root of the diagonal elements of the variance covariance matrix. The replication code for obtaining asymptotic standard error used throughout this paper is provided in the supplementary material.

The magnitude of change is reflected in d . When d is 0, it represents homogeneity in the corresponding parameter, which serves as the baseline; when d is greater than 0, it represents heterogeneity in V (increasing with Cognitive Ability in this example), with larger d indicating more severe heterogeneity. One example of data with and without heterogeneity is displayed in Fig. 1. In the left panel, data were generated without heterogeneity in a random slope ($d=0$), whereas in the right panel data were generated with heterogeneity in a random slope as large as $d=4$. Within each panel, the subjects 1–6 denoted with gray lines have cognitive ability equal to or less than 2, whereas subjects 7–12 denoted with yellow lines have cognitive ability greater than 2. Without the impact of variance heterogeneity (left panel), it is easy to observe that the RT has a positive relation with Days (β_1), and this relation differs for subjects with different cognitive abilities (β_3). However, these relations are difficult to see under the impact of variance change (right panel). Unfortunately, the real data often look more similar to the right panel, with no obvious relations to be detected even if the generating fixed parameters are actually exactly the same as

the left panel. To formally examine the impact of heterogeneity, we computed the percentage of significant fixed parameters related to “Days” (β_1 and β_3) among 1000 replications in each condition.

The full simulation results for β_1 and β_3 are demonstrated in Fig. 2, with the panel titles first indicating the tested parameter and then indicating the heterogeneous parameter, and the y-axis representing power (using $\alpha = 0.05$). In general, when sample size is medium or large, increasing heterogeneity in the slope variance σ_1^2 or covariance σ_{01} reduces power for both the main effect and interaction effect. Heterogeneity in the residual variance or intercept variance does not impact power for β_1 or β_3 , because they can be compensated for during estimation (Kwok et al., 2007). That is to say, when the intercept variance (or residual variance) increases, the residual variance (or intercept variance) estimate will decrease to compensate for the change, leading to the diagonal of V being unchanged. This compensation effect exists because the intercept covariate in the random effect design matrix (Z) is all 1, so that the intercept and residual variance contribute equally to the diagonal of V .

When sample size is small ($n=120$), power is generally lower in all scenarios. In addition, greater heterogeneity in the residual variance also leads to lower power, which might be due to the fact that heterogeneity combined with small sample size is more likely to result in unstable variance/covariance estimates, or even convergence issues. Overall, however, failing to account for the upward changes in V would generally result in type II error.

Although it is important to systematically monitor heterogeneity in variance components, it is also plausible that a fixed parameter indeed changes according to another variable (e.g., that an interaction exists). Ideally, there would exist a

Fig. 1 Simulated data sets without heterogeneity (left panel) and with heterogeneity (right panel). The sample size for both examples is 120.

Fig. 2 Simulated power curves for β_1 and β_3 under situations with heterogeneity in parameters σ_r^2 , σ_1^2 , σ_0^2 , and σ_{01} across ranges of 0–4 asymptotic variance. Panel labels denote the parameter being tested and the parameter with heterogeneity.

statistical test that can differentiate between these two kinds of changes. In the next section, we will introduce a score-based family of statistical tests that can fulfill this need.

Score-based tests

In this section, we will introduce the score-based test as applied to the framework of LMM. This introduction draws on LMM results described by Wang & Merkle (2018) and is related to tests described by, e.g., Zeileis & Hornik (2007), Merkle et al. (2014), and Wang et al. (2014).

Scores

Based on the marginal model expression shown in Eq. 13, the log likelihood of the LMM can be expressed as:

$$\ell(\sigma^2, \beta; \mathbf{y}) = -\frac{n}{2} \log(2\pi) - \frac{1}{2} \log(|\mathbf{V}|) - \frac{1}{2} (\mathbf{y} - \mathbf{X}\beta)^\top \mathbf{V}^{-1} (\mathbf{y} - \mathbf{X}\beta). \quad (20)$$

Scores, denoted $s_i()$ in this paper, are based on the first partial derivatives of ℓ w.r.t. $\theta = (\sigma^2, \beta)^\top$. The scores involve these partial derivatives evaluated for each observation i , where $i = 1, 2, \dots, n$, and they can be roughly viewed as a residual: values close to 0 imply that the model provides a good fit to case i with respect to a specific parameter, and values far from 0 imply the opposite.

The model gradient is equal to the sum of scores across all individuals and clusters:

$$\ell'(\theta; \mathbf{y}) = \sum_{j=1}^J \frac{\partial \ell(\theta; \mathbf{y}_j)}{\partial \theta} = \sum_{j=1}^J \sum_{i \in c_j} s_i(\theta; \mathbf{y}_i) \quad (21)$$

where $\frac{\partial \ell(\theta; \mathbf{y}_j)}{\partial \theta}$ represents the first derivative within cluster c_j , which can be expressed as the sum of the case-wise score $s_i()$ belonging to cluster j . For LMMs, the function $s_i()$ w.r.t. σ_k^2 and β can be expressed as the i th component of the vectors or as the i th row of the matrix (if β contains multiple components):

$$s(\sigma_k^2; \mathbf{y}) = -\frac{1}{2} \text{diag} \left[\mathbf{V}^{-1} \frac{\partial \mathbf{V}}{\partial \sigma_k^2} \right] + \left\{ \frac{1}{2} (\mathbf{y} - \mathbf{X}\beta)^\top \mathbf{V}^{-1} \left(\frac{\partial \mathbf{V}}{\partial \sigma_k^2} \right) \mathbf{V}^{-1} \right\}^\top \cdot (\mathbf{y} - \mathbf{X}\beta) \quad (22)$$

$$s(\beta; \mathbf{y}) = \{ \mathbf{X}^\top \mathbf{V}^{-1} \}^\top \cdot (\mathbf{y} - \mathbf{X}\beta), \quad (23)$$

where $\cdot$ represents the Hadamard product (component-wise multiplication). Further detail on these derivations can be found in Wang & Merkle (2018), McCulloch & Neuhaus (2001), and Stroup (2012).

These equations provide scores for each observation i , and we can construct the cluster-wise scores by summing scores within each cluster. In situations with one grouping (clustering) variable, the cluster-wise scores can be obtained from a fitted *lme4* model via R package *merDeriv* (Wang & Merkle, 2017).

Statistics

As applied to the LMMs considered here, score-based tests can be used to study heterogeneity that is potentially explained by

an auxiliary variable T ; for example, in the data demonstration considered earlier, the auxiliary variable could have been Cognitive Ability. Because the scores can be viewed as a type of residual, the score-based tests basically help us judge whether the residual magnitudes are associated with T . Because we have unique scores for each model parameter, we can also obtain information about where heterogeneity occurs.

Statistically, the tests considered here can be viewed as generalizations of the Lagrange multiplier test. The tests are based on a cumulative sum of scores, where the order of accumulation is determined by T . If there is no heterogeneity explained by T , then this cumulative sum should fluctuate around zero. Otherwise, the cumulative sum would systematically diverge from zero.

To formalize these ideas, we first define the (scaled) cumulative sum of the ordered scores. This can be written as

$$\mathbf{B}(t; \hat{\boldsymbol{\theta}}) = \hat{\mathbf{I}}^{-1/2} \mathcal{J}^{-1/2} \sum_{j=1}^{\lfloor Jt \rfloor} s(\hat{\boldsymbol{\theta}}; y_{(j)}) \quad (0 \leq t \leq 1) \quad (24)$$

where $\hat{\mathbf{I}}$ is an estimate of the information matrix, $\lfloor Jt \rfloor$ is the integer part of Jt (i.e., a floor operator), and $x_{(j)}$ reflects the cluster with the j th smallest value of the auxiliary variable T . While the above equation is written in general form, we can restrict the value of t in finite samples to the set $\{0, 1/J, 2/J, 3/J, \dots, J/J\}$. We focus on how the cumulative sum fluctuates as more clusters' scores are added to it, e.g., starting with the person of lowest cognitive ability and ending with the person of highest cognitive ability. The summation is pre-multiplied by an estimate of the inverse square root of the information matrix, which serves to decorrelate the fluctuation processes associated with model parameters. For LMMs, $\hat{\mathbf{I}}$ can be written as expected information matrix (e.g., Wang & Merkle, 2018):

$$\hat{\mathbf{I}} = \begin{bmatrix} \mathbf{X}^T \mathbf{V}^{-1} \mathbf{X} & 0 \\ 0 & \left(\frac{1}{2}\right) \text{tr} \left[\mathbf{V}^{-1} \left(\frac{\partial \mathbf{V}}{\partial \sigma_{k1}^2} \right) \mathbf{V}^{-1} \left(\frac{\partial \mathbf{V}}{\partial \sigma_{k2}^2} \right) \right] \end{bmatrix}.$$

Score-based tests statistics

To obtain an official test statistic, we must summarize the behavior of the cumulative sum in a scalar. Multiple summaries are available, leading to multiple tests of the same hypothesis. For example, one could take the absolute maximum that the cumulative sum attains for any parameter of interest, resulting in a *double max* statistic (the maximum is taken across parameters and clusters entering into the cumulative sum). Alternatively, one could sum the (squared) cumulative sum across parameters of interest and take the maximum or

the average across clusters, resulting in a *maximum Lagrange multiplier* statistic and *Cramér-von Mises* statistic, respectively (see Merkle & Zeileis, 2013 for further discussion). These statistics are given by

$$DM = \max_{j=1, \dots, J} \max_{k=1, \dots, K} |\mathbf{B}(\hat{\boldsymbol{\theta}})_{jk}| \quad (25)$$

$$CvM = \mathcal{J}^{-1} \sum_{j=1, \dots, J} \sum_{k=1, \dots, K} \mathbf{B}(\hat{\boldsymbol{\theta}})_{jk}^2, \quad (26)$$

$$\max LM = \max_{j=j, \dots, J} \left\{ \frac{j}{J} \left(1 - \frac{j}{J} \right) \right\}^{-1} \sum_{k=1, \dots, K} \mathbf{B}(\hat{\boldsymbol{\theta}})_{jk}^2. \quad (27)$$

For an ordinal auxiliary variable T with m levels, we can modify the statistics above so that the maximum is only considered after all clusters at the same level of T have entered the summation. This leads to test statistics that are especially sensitive to heterogeneity that is monotonic with T (Merkle et al., 2014). Formally, we define t_L ($L = 1, \dots, m-1$) to be the empirical, cumulative proportions of clusters observed at the first $m-L$ levels of T . The modified statistics are then given by

$$WDM_o = \max_{j \in \{j_1, \dots, j_{m-1}\}} \left\{ \frac{j}{J} \left(1 - \frac{j}{J} \right) \right\}^{-1/2} \max_{k=1, \dots, K} |\mathbf{B}(\hat{\boldsymbol{\theta}})_{jk}|, \quad (28)$$

$$\max LM_o = \max_{j \in \{j_1, \dots, j_{m-1}\}} \left\{ \frac{j}{J} \left(1 - \frac{j}{J} \right) \right\}^{-1} \sum_{k=1, \dots, K} \mathbf{B}(\hat{\boldsymbol{\theta}})_{jk}^2, \quad (29)$$

where $j_L = \lfloor n \cdot t_L \rfloor$ ($L = 1, \dots, m-1$). Finally, if the auxiliary variable T is only nominal/categorical, the cumulative sums of scores can be used to obtain a Lagrange multiplier statistic. This test statistic can be formally written as

$$LM_{uo} = \sum_{L=1, \dots, m} \sum_{k=1, \dots, K} \left(\mathbf{B}(\hat{\boldsymbol{\theta}})_{j_L k} - \mathbf{B}(\hat{\boldsymbol{\theta}})_{j_{L-1} k} \right)^2, \quad (30)$$

where $\mathbf{B}(\hat{\boldsymbol{\theta}})_{j_0 k} = 0$ for all k . This statistic is asymptotically equivalent to the usual, likelihood ratio test statistic, and it is advantageous over the likelihood ratio test because it requires estimation of only one model (the restricted model). We make use of this advantage in the simulations, described later.

In the following section, we apply these tests to a linear mixed model with one grouping variable, studying the tests' ability to distinguish between heterogeneity and interactions.

Simulation

The goal of the simulation is to examine score-based tests' abilities to differentiate between changes in fixed effect parameters (i.e., interaction effects) and changes in variance parameters (i.e., heterogeneity). For ease of description, we frame the data-generating model as being based on a longitudinal depression intervention administered to

participants with different levels of cognitive ability (here we assume that $m = 4$, i.e., that there are four ordered levels of cognitive ability). Each participant's depression magnitude is measured once per month during a 10-month period. Thus 10 measurements are nested within each participant, which comprises a typical application for LMMs. It is plausible that the amount of time needed to change the magnitude of depression is dependent on subjects' cognitive ability. If so, there exists an interaction between time and cognitive ability. However, it is also possible that patients with higher cognitive ability have larger intercept variance (σ_0^2) or residual variance (σ_r^2). In addition, the interaction and heterogeneity might occur simultaneously. Since both interaction and heterogeneity can be viewed as parameter instability w.r.t. an auxiliary variable, we aim to examine the extent to which the score-based tests could attribute the parameter instability to the truly changing parameter(s) in an LMM framework.

Method

Data were generated from an LMM. The predictor is time, with its associated coefficient as β_1 , and β_0 serves as the fixed intercept, which completes the fixed parameters in the model. We have covarying intercept and slope random effects as well, with the variance and covariance as σ_0^2 , σ_{01} , and σ_1^2 . The variance not captured by the random effects is modeled by the residual variance σ_r^2 . The true parameter change can occur in one of seven ways: fixed intercept β_0 , time coefficient β_1 , random intercept variance σ_0^2 , random covariance σ_{01} , random slope variance σ_1^2 , residual variance σ_r^2 , or β_1 and σ_r^2 simultaneously. The fitted models matched the data-generating model, and parameter estimates were obtained by marginal maximum likelihood. Parameter changes were tested in each of the six estimated parameters.

Power and type I error were examined across three sample sizes ($n = 120, 480, 960$), five magnitudes of parameter change, and six tests of each individual parameter. The parameters' true values were set to be the same as the estimates from the *sleepstudy* data included in *lme4*. The parameter change point and changing magnitude is manipulated in the same way as the problem demonstration simulation.

For each combination of sample size (n) $\times$ violation magnitude (d), 1000 data sets were generated, and all parameters were tested. Two ordinal statistics ($\max LM_o$, WDM_o) and one categorical statistic (LM_{uo}) were examined (Merkle et al., 2014; Wang et al., 2014). The categorical statistic is asymptotically equivalent to the usual likelihood ratio test. Thus, this statistic provides information about the relative performance of the ordinal statistics vs. the LRT.

Results

Full simulation results are presented in Figs. 3, 4, 5, 6, 7, 8, and 9. In each graph, the x-axis represents the violation magnitude and the y-axis represents the power of detecting parameter change. When $x = 0$, the corresponding power serves as the type I error. Fig. 3 demonstrates power curves as a function of violation magnitude in β_0 , with sample size changing across rows, the tested parameters changing across columns, and lines reflecting different test statistics. Figs. 4, 5, 6, 7, and 8 display similar power curves when the true changing parameter is β_1 , σ_0 , σ_{01} , σ_1 , and σ_r^2 , respectively. Figure 9 shows the power curves when there exist two changing parameters, β_1 and σ_r^2 .

From these figures, one can generally observe that the score-based statistics could isolate the truly changing parameter, with non-zero power curves for changing parameter(s), and near-zero power curves for non-changing parameters. For example, in Fig. 3, for β_0 , the power increases with the violation magnitude d and sample size (across rows); the power for the other five non-changing parameters remain near zero (across columns), even with increasing violation magnitude and sample size.

Within each non-zero power curve panel of Figs. 3, 4, 5, 6, 7, 8, and 9, the two ordinal statistics, $\max LM_o$ and WDM_o , exhibit higher (when testing fixed parameter or random intercept variance) or similar (when testing residual variance) power compared with categorical statistic LM_{uo} . This is partially consistent with the results demonstrated in Merkle et al., (2014), where ordinal statistics are shown to be more sensitive to monotonic parameter changes. The residual variance results (Fig. 8) might be due to a ceiling effect, where all three power curves quickly increase to 1. In conditions with only one changing parameter (Figs. 3, 4, 5, 6, 7, and 8), $\max LM_o$ and WDM_o are mathematically equivalent (Merkle et al., 2014). In conditions with two changing parameters (Fig. 9), $\max LM_o$ and WDM_o still demonstrate similar power curves. The advantages of WDM_o are only apparent when testing many (more than two) parameters at a time Merkle et al. (2014); Wang et al. (2014).

Comparing the non-zero power curves across these seven figures, it shows the score-based tests have somewhat higher power to detect residual variance change when sample size is medium or large, followed by fixed parameter change and random variance/covariance parameter change. This phenomenon is most apparent by comparing Figs. 5, 6, and 7 with Fig. 9, with the power curve for the residual variance and fixed parameter approaching 1 in conditions with medium or large sample sizes, while the power curve for the random variance/covariance ranges from 0.4 to 0.8 even for the greatest d under large sample size. The general difficulty to detect parameter changes in the G matrix is related to the fact that large

Fig. 3 Simulated power curves for max LM_o , WDM_o , and LM_{uo} across parameter change of 0–4 asymptotic standard error. The changing parameter is β_0 . Panel labels denote the parameter being tested along with sample size.

parameter changes in variance/covariance components often render $\mathbf{G}$ as numerically non-positive definite, resulting in correlations of 1 or model non-convergence. The Discussion section provides more details on this issue.

In summary, we found that the score-based tests can attribute heterogeneity to the truly problematic parameter(s) in an LMM context. Additionally, the tests were more sensitive to changes in fixed effect parameters, as

Fig. 4 Simulated power curves for max LM_o , WDM_o , and LM_{uo} across parameter change of 0–4 asymptotic standard error. The changing parameter is β_1 . Panel labels denote the parameter being tested along with sample size.

Fig. 5 Simulated power curves for max LM_o , WDM_o , and LM_{uo} across parameter change of 0–4 asymptotic standard error. The changing parameter is σ_0^2 . Panel labels denote the parameter being tested along with sample size.

compared to variance parameters. In the next section, we apply the tests to real data to illustrate the potential usage of score-based tests in practice. The general

approach is to fit an LMM of interest, then obtain each parameter's score-based test statistics w.r.t. an auxiliary variable in level 2. If the variance (either random effect or

Fig. 6 Simulated power curves for max LM_o , WDM_o , and LM_{uo} across parameter change of 0–4 asymptotic standard error. The changing parameter is σ_{01} . Panel labels denote the parameter being tested along with sample size.

Fig. 7 Simulated power curves for max LM_o , WDM_o , and LM_{uo} across parameter change of 0–4 asymptotic standard error. The changing parameter is σ_1^2 . Panel labels denote the parameter being tested along with sample size.

residual) component is detected to have parameter instability, it indicates heterogeneity present in the data; if the fixed parameter demonstrates instability, then we can claim interaction between the covariate and the auxiliary variable.

Tutorial

In this section, we demonstrate how the score-based tests can be carried out in R, using package *lme4* (Bates, Mächler, Bolker, &

Fig. 8 Simulated power curves for max LM_o , WDM_o , and LM_{uo} across parameter change of 0–4 asymptotic standard error. The changing parameter is σ_r^2 . Panel labels denote the parameter being tested along with sample size.

Fig. 9 Simulated power curves for max LM_o , WDM_o , and LM_{uo} across parameter change of 0–4 asymptotic standard error. The changing parameter are β_1 and σ_r^2 . Panel labels denote the parameter being tested along with sample size.

Walker, 2015) for model fit and *strucchange* (Zeileis, Leisch, Hornik, & Kleiber, 2002; Zeileis, 2006) for testing, with the score computations handled in the background by *merDeriv* (Wang & Merkle, 2017). We use data from the 1982 “High School and Beyond” survey funded by the National Center for Education Statistics (NCES), which is available in the R package *mlmRev*. In the dataset, 7185 U.S. high-school students from 160 schools completed a math achievement test, with the students’ socioeconomic status (*ses*) as a level 1 covariate. We center the *ses* covariate by school mean and focus on the centered *ses* (denoted as *cses*) below, as recommended by previous researches (Algina & Swaminathan, 2011; Bauer & Curran, 2005; Enders & Tofghi, 2007). The centering only eases parameter interpretation, and generally has no impact on the cross-level interaction term’s statistical significance test (Algina & Swaminathan, 2011).

The aim of the current analyses is to determine how students’ math achievement scores (denoted as *mAch* in the dataset) are associated with their family socioeconomic status. It is plausible that the relationship between *cses* and math achievement differs across schools with different *meanses* (level 2 covariate). The traditional approach is to fit the linear mixed model with an interaction term and examine the significance of the coefficient for the interaction term.

Traditional approach

The traditional approach to testing the interaction between *cses* and *meanses* can be carried out via.

```
> library("mlmRev")
> library("lme4")
> library("lmerTest")
> m2 <- lmer(mAch ~ cses*meanses +
(cses|school), data = Hsb82, REML = FALSE)
> summary(m2)
```

where *cses*meanses* specifies the model fixed effects (both main effects and interaction), and (*cses|schools*) specifies the random effects; *REML = FALSE* requests the marginal maximum likelihood estimation as described in Eq. 20. From the results returned by *summary()* (not shown), we can see the coefficient for the interaction term is not significant ($p = 0.367$). However, as shown in the Fig. 2, the significance test for the interaction might be impacted by variance/covariance heterogeneity in random effects (second row of Fig. 2, second panel and fourth panel). Thus, we use the score-based tests to distinguish between the cross-level interaction and variance heterogeneity.

To conduct the score-based tests, we first fit the model with only the level 1 covariate, as shown in the code section below. One advantage of score-based tests is that the focal level 2 covariate does not enter the model but serves as the auxiliary variable in the testing stage. This feature reduces model complexity and is more likely to lead to converged models.

```
> library("mlmRev")
> library("lme4")
> m1 <- lmer(mAch ~ cses + (cses | school),
data = Hsb82, REML = FALSE)
```

This fitted model includes six parameters (with labels 1–6 below), which are β_0 , β_1 , σ_0^2 , σ_{01} , σ_1^2 , and σ_r^2 (in this order). While the second-level covariate `meanses` would generally be treated as continuous; for demonstration purposes, we consider treating it as continuous, ordinal, and categorical. Each of these treatments is described separately below.

Continuous treatment

If we treat the auxiliary variable `meanses` as continuous, we can employ `sctest()` to obtain continuous statistics from Eqs. 25 to 27 for the parameter of interest, which is specified by the `parm` argument.

Because `sctest()` utilizes `estfun()` in the background, we need to ensure that the ordering of the auxiliary variable `meanses` matches the row ordering of the `estfun()` output, so that each value of `meanses` corresponds to its associated `school`. The ordering and checking can be completed by the code below. This step is highly recommended in practice; the *data.table* package (Dowle et al., 2019) is utilized here solely for speed purposes.

```
> ## data.table for speed.
> library("data.table")
> Hsb82 <- as.data.table(Hsb82)
> ## ordering
> orderHsb82 <- Hsb82[order(as.numeric(
school))]]
> ## checking
> all.equal(rownames(estfun(m1)), unique
(levels(orderHsb82$school)))
```

After the ordering checking, we can proceed to the statistical tests. For example, we can test whether the random slope variance (σ_1 , specified by `parm = 5`) is stable across `meanses`. The code below displays how to conduct the tests.

```
> library("strucchange")
> library("merDeriv")
> dm <- sctest(m1, order.by = unique(orderHsb82$meanses),
parm = 5, functional = "DM")
> cvm <- sctest(m1, order.by = unique(orderHsb82$meanses),
parm = 5, functional = "CvM")
> maxlm <- sctest(m1, order.by = unique(orderHsb82$meanses),
parm = 5, functional = "maxLM")
> c(dm$p.value, cvm$p.value, maxlm$p.value)
```

[1] 0.013039947 0.009432531 0.002934033

We can see that all three statistics indicate significant parameter instability for the random slope variance parameter, suggesting the existence of heterogeneity.

Ordinal treatment

In some scenarios, the school-level `ses` variable may only be measurable as ordered categories. To mimic this situation, we categorize schools with similar `meanses` to yield an ordinal level 2 auxiliary variable with five categories. The code below first creates an ordinal variable and then shows that the only change in the `sctest()` command is the `functional` argument.

```
> # create ordinal variable
> orderHsb82[meanses < -0.5, schses:=1]
> orderHsb82[meanses <= -0.1 & meanses >
-0.5, schses:=2]
> orderHsb82[meanses <= 0.1 & meanses >
-0.1, schses:=3]
> orderHsb82[meanses <= 0.45 & meanses >
0.1, schses:=4]
> orderHsb82[meanses >= 0.45, schses:=5]
> # compute test statistics
> ordervar <- orderHsb82[,mean(as.numeric(schses)),
by=school]$V1
> wdm <- sctest(m1, order.by = ordervar, parm = 5,
+ functional = "WDMo")
> maxlmo <- sctest(m1, order.by = ordervar, parm = 5,
+ functional = "maxLMo")
> c(wdm$p.value, maxlmo$p.value)
```

[1] 0.02977961 0.03029124

Like the previous statistics, both statistics are significant here as well.

Categorical treatment

Lastly, when there is no ordering information contained in the auxiliary variable, a categorical statistic can be implemented in the following way. This statistic is asymptotically equivalent to the traditional LRT as stated before, but has less power of detecting change as demonstrated in the simulation. In this example, the test result is not significant ($\alpha = 0.05$) because the ordering information was ignored.

```
> lmuo <- sctest(m1, order.by = ordervar,
parm = 5, functional = "LMuo")
> lmuo$p.value
```

[1] 0.06175341

Subgroup information

In addition to the test statistics and p value, “instability plots” can be generated by setting `plot = TRUE` in the `sctest()` functions above. Figure 10 displays the ordinal statistics’ fluctuation across `schses` levels. In this figure, the first column displays the fluctuation process associated with $\max LM_o$, and the second column displays the fluctuation process associated with WDM_o . Each panel represents the test of a specific model parameter, shown in the panel title. Within each panel, the horizontal dashed line represents the 5% critical value. If the solid line crosses the critical value, then there is evidence that the corresponding parameter fluctuates across `schses` (because the full set of scores sum to zero, the final level of `schses` is not displayed on the x-axis).

In Fig. 10, it is observed that the β_0 , β_1 , σ_0^2 , and σ_1^2 demonstrate parameter instability, whereas σ_{01} and σ_r^2 do not. The instability of β_0 indicates that there exists a main effect of `schses`, and the instability of β_1 implies that there exists a cross-level interaction effect between `schses` and `cses`. In addition, the random intercept and the random slope demonstrate instability. As described earlier, this heterogeneity in random effect variances appears to have “masked” the significance of the interaction term.

Figure 10 also provides information about levels of `schses` where parameters differ from one another; this can be discerned from levels where the solid line crosses the dashed horizontal line. Thus, the intercept parameter changes w.r.t. each of the four levels of `schses` (all points are above the line); the slope of `ses` (β_1) at `schses` level 1 differs from the other levels; the random intercept variance σ_0^2 has two changing points: one at level 1 and the other at level 4; and the random slope σ_1^2 differs between level 1 and the remaining levels. These results provide more detailed information about how `schses` is associated with different pieces of the model.

To expand on the results above, we can create a dummy variable for students at `schses` level 1 (coded as 0), then use that dummy variable in place of `meanses`. The code is given below:

```
> Hsb82[schses==1, dummy := 0]
> Hsb82[schses!=1, dummy := 1]
> m3 <- lmer(mAch ~ cses*dummy + (cses | school),
data = Hsb82, REML = FALSE)
```

The interaction between `cses` and the dummy variable is 1.09 with $p < 0.05$, indicating a significant interaction. Thus, informed by the instability plots, we can detect the “masked” interaction effect via traditional methods. Alternatively, we can fit the model separately for students at schools with the lowest level of `schses` and for students at other schools with larger values of `schses`. The former model results in β_1 (coefficient of `cses`) as 1.21, whereas the latter model has

β_1 as 2.31. These results indicate that students’ `cses` has stronger relationship with math achievement in schools with higher `ses`.

In summary, score-based tests provide a statistical tool to closely examine an LMM’s parameter estimates with respect to an auxiliary, level 2 variable. The examination of variance components (random effect (co) variance and residual variance) provides tests of heterogeneity. Additionally, the fluctuation plots can be used to interpret the nature of heterogeneity or interactions, without arbitrary median splits or subsamples of data with few observations.

Discussion

In this paper, we extended a family of score-based tests to linear mixed models, focusing on models with one grouping variable. We found that the tests can isolate specific parameters that exhibit instability, which avoids masked cross-level interaction effects in the presence of heterogeneity. They also provide specific information about groupings of the auxiliary variable whose parameter values differ. The tests developed in this paper can currently only be carried out on an auxiliary variable measured at the model’s upper level (level 2), a restriction that leads to the future directions described below.

Grouping with multiple variables

The auxiliary variable is specifically required to be at the upper level because the tests described here require that the scores be independent. This independence assumption challenges models with at least two variables defining clusters, such as models with (partially) crossed random effects, or models with multilevel nested designs (e.g., Bates, 2010, Ch. 2). In these cases, we cannot simply sum scores within a cluster to obtain independent, cluster-wise scores, because observations in different clusters on the first grouping variable may be in the same cluster on the second grouping variable. A related issue occurs when the auxiliary variable is at the lowest (first) level of the model: scores at the lowest level are not independent, so the tests described here cannot be immediately used to test parameters with respect to a level 1 variable.

A natural approach to deal with the issue of dependent scores is to find a heteroskedasticity and autocorrelation consistent (HAC) covariance matrix estimator. The traditional Hessian matrix only accounts for the correlations among score columns, whereas the HAC estimator is a robust Hessian estimator that can serve the purpose of de-correlating the scores under a generalized linear model framework. Several methods have been proposed here, including kernel HAC estimators with automatic bandwidth selection and weighted empirical adaptive

Fig. 10 Empirical fluctuation processes of the $\max LM_o$ statistic (first column) and WDM_o (second column) for β_0 (first row), β_1 (second row), σ_0^2 (third row), σ_{01} (fourth row), σ_1^2 (fifth row) and σ_r^2 (sixth row),

using M1 model. The statistics corresponding to the 5th numeracy level within each panel always equal to 0, so not shown in the panels.

variance estimators (Andrews, 1993). For multilevel models with multiple grouping variables, Berger, Graham, & Zeileis (2017) recently implemented a sandwich approach to obtaining robust variance/covariance matrices. It might be possible to deploy these methods in the context of linear mixed models.

In practice, however, the technical challenge is to find the optimal bandwidth in empirical studies with non-independent observations. Suboptimal selection of the bandwidth parameter could lead to decreasing power of detecting a parameter change or even drop to zero (Perron, 2006). Shao & Zhang (2010) recently proposed use of a “self-normalizing” approach to tackle this technical issue, and they have utilized this approach in multivariate settings (Zhang, Shao, Hayhoe, Wuebbles, et al., 2011). An extension to linear mixed models with non-independent level 2 grouping variables and non-diagonal residual covariance matrices is currently under development.

Model estimation

Along with independence issues, general model estimation issues may influence the score-based tests’ accuracies. For example, in the relatively common case where a parameter estimate lies on the boundary (e.g., a correlation between random effects near ± 1 or a variance approaching zero), then it may be impossible to carry out the proposed tests due to the non-positive definite structure of the model information matrix. Additionally, model misspecification can also influence the tests. One common type of misspecification involves the residual covariance matrix having nonzero, off-diagonal elements that are fixed to zero in the estimated model. Wang et al. (2014) examined the tests’ performance in the factor analysis framework, and they found that the instability of unmodeled parameters would lead the tests to identify instability in related model parameters. In the same manner, we speculate that instability in unmodeled, off-diagonal residual covariances would be incorrectly attributed to random effects’ variances or covariances (**G** components). Thus, it is important to carefully consider the specification of the estimated model.

Tests’ power

The power curves demonstrated in the simulation section are related to the parameters’ asymptotic standard errors and to sample size. In the simulation, we used the asymptotic variance-covariance matrix from the `sleepstudy` data, scaling these asymptotic standard errors by the square root of the sample size. However, as demonstrated in the tutorial section, applied researchers do not need to obtain the asymptotic standard error to utilize score-based tests. Further, as demonstrated in the simulation, the power generally increases with sample size. The simulation sample sizes of 120, 480, 960 were used to conveniently allocate observations in four levels. To achieve high power, there is a trade-off between

sample size needed and the magnitude of instability. In the simulation, the small sample size of 120 was sufficient because the instability was large enough.

Summary

In this paper, we generalized a family of score-based tests to two-level linear mixed models, which allow researchers to test whether model parameters fluctuate with an unmodeled level 2 variable. We found that the tests could successfully decouple cross-level interactions from variance heterogeneity, whereas heterogeneity could cause the traditional significance test of a cross-level interaction to exhibit inflated type II error. Along with providing information about parameter stability across all estimated LMM parameters, the tests provide additional information about heterogeneous subgroups when parameter instability is detected. Thus, applied researchers in psychology and education can use the tests to examine potential cross-level interactions while ruling out possible masked results due to heterogeneity.

Computational Details All results were obtained using the R system for statistical computing (R Core Team, 2018), version 3.6.1, employing the add-on package *lme4* 1.1–21 (Bates, Mächler, et al., 2015) for fitting of the linear mixed models and *strucchange* 1.5–2 (Zeileis et al., 2002; Zeileis, 2006) for evaluating the parameter instability tests. R and both packages are freely available under the General Public License from the Comprehensive R Archive Network at <http://CRAN.R-project.org/>. R code for replication of our results is available at <http://semtools.R-Forge.R-project.org/>.

Acknowledgements This work was supported by National Science Foundation grants SES-1061334 and 1460719.

References

- Abdolell, M., LeBlanc, M., Stephens, D., and Harrison, R. (2002). Binary partitioning for continuous longitudinal data: categorizing a prognostic variable. *Statistics in medicine*, 21(22):3395–3409.
- Abou El-Komboz, B., Zeileis, A., and Strobl, C. (2014). Detecting differential item and step functioning with rating scale and partial credit trees. Technical Report 152, Department of Statistics, Ludwig-Maximilians-Universität München.
- Algina, J. and Swaminathan, H. (2011). Centering in two-level nested designs. *Handbook of advanced multilevel analysis*, pages 285–312.
- Andrews, D. W. K. (1993). Tests for parameter instability and structural change with unknown change point. *Econometrica*, 61:821–856.
- Bates, D. (2010). *lme4: Mixed-effects Modeling with R*. New York: Springer.
- Bates, D., Kliegl, R., Vasishth, S., and Baayen, H. (2015). Parsimonious mixed models. *arXiv preprint arXiv:1506.04967*.

- Bates, D., Mächler, M., Bolker, B., and Walker, S. (2015). Fitting linear mixed-effects models using *lme4*. *Journal of Statistical Software*, 67(1):1–48.
- Bauer, D. J. and Curran, P. J. (2005). Probing interactions in fixed and multilevel regression: Inferential and graphical techniques. *Multivariate behavioral research*, 40(3):373–400.
- Belenky, G., Wendenstein, N. J., Thorne, D. R., Thomas, M. L., Sing, H. C., Redmond, D. P., Russo, M. B., and Balkin, T. J. (2003). Patterns of performance degradation and restoration during sleep restriction and subsequent recovery: A sleep dose-response study. *Journal of Sleep Research*, 12(1):1–12.
- Berger, S., Graham, N., and Zeileis, A. (2017). Various versatile variances: An object-oriented implementation of clustered covariances in R. Technical report.
- Brandmaier, A. M., von Oertzen, T., McArdle, J. J., and Lindenberger, U. (2013). Structural equation model trees. *Psychological methods*, 18(1):71.
- Bresin, K. (2014). Five indices of emotion regulation in participants with a history of nonsuicidal self-injury: a daily diary study. *Behavior Therapy*, 45(1):56–66.
- R Core Team (2018). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Dowle, M., Srinivasan, A., Gorecki, J., Chirico, M., Stetsenko, P., Short, T., Lianoglou, S., Antonyan, E., Bonsch, M., Parsonage, H., et al. (2019). *Package “data.table”: Extension of data.frame*. R package version 1.12.8.
- Ebner-Priemer, U. W., Houben, M., Santangelo, P., Kleindienst, N., Tuerlinckx, F., Oravecz, Z., Verleysen, G., Van Deun, K., Bohus, M., and Kuppens, P. (2015). Unraveling affective dysregulation in borderline personality disorder: A theoretical model and empirical evidence. *Journal of Abnormal Psychology*, 124(1):186.
- Enders, C. K. and Tofighi, D. (2007). Centering predictor variables in cross-sectional multilevel models: a new look at an old issue. *Psychological methods*, 12(2):121.
- Fokkema, M., Smits, N., Zeileis, A., Hothorn, T., and Kelderman, H. (2018). Detecting treatment-subgroup interactions in clustered data with generalized linear mixed-effects model trees. *Behavior Research Methods*, 50(5):2016–2034.
- Halaby, C. N. (2004). Panel models in sociological research: Theory into practice. *Annu Rev Sociol.*, 30:507–544.
- Knight, A. P. (2013). Mood at the midpoint: Affect and change in exploratory search over time in teams that face a deadline. *Organization Science*, 26(1):99–118.
- Koval, P. and Kuppens, P. (2012). Changing emotion dynamics: Individual differences in the effect of anticipatory social stress on emotional inertia. *Emotion*, 12(2):256.
- Kwok, O.-M., West, S. G., and Green, S. B. (2007). The impact of misspecifying the within-subject covariance structure in multiwave longitudinal multilevel models: A Monte Carlo study. *Multivariate Behavioral Research*, 42(3):557–592.
- Leckie, G., French, R., Charlton, C., and Browne, W. (2014). Modeling heterogeneous variance–covariance components in two-level models. *Journal of Educational and Behavioral Statistics*, 39(5):307–332.
- Lockwood, J., McCaffrey, D. F., et al. (2007). Controlling for individual heterogeneity in longitudinal models, with applications to student achievement. *Electronic Journal of Statistics*, 1:223–252.
- McCulloch, C. E. and Neuhaus, J. M. (2001). *Generalized Linear Mixed Models*. New York: John Wiley & Sons.
- McLachlan, G. J. and Basford, K. E. (1988). *Mixture models: Inference and applications to clustering*, volume 84. Marcel Dekker.
- Merkle, E. C., Fan, J., and Zeileis, A. (2014). Testing for measurement invariance with respect to an ordinal variable. *Psychometrika*, 79:569–584.
- Merkle, E. C. and Zeileis, A. (2013). Tests of measurement invariance without subgroups: A generalization of classical methods. *Psychometrika*, 78:59–82.
- Perron, P. (2006). Dealing with structural breaks. In *Palgrave handbook of econometrics*, volume 1, pages 278–352.
- Röcke, C., Li, S.-C., and Smith, J. (2009). Intraindividual variability in positive and negative affect over 45 days: Do older adults fluctuate less than young adults? *Psychology and Aging*, 24(4):863.
- Shao, X. and Zhang, X. (2010). Testing for change points in time series. *Journal of the American Statistical Association*, 105(491):1228–1240.
- Stegmann, G., Jacobucci, R., Serang, S., and Grimm, K. J. (2018). Recursive partitioning with nonlinear models of change. *Multivariate behavioral research*, 53(4):559–570.
- Stephens, M. A. (1974). EDF statistics for goodness of fit and some comparisons. *Journal of the American statistical Association*, 69(347):730–737.
- Stroup, W. W. (2012). *Generalized Linear Mixed Models: Modern Concepts, Methods and Applications*. New York: CRC Press.
- Thompson, R. J., Mata, J., Jaeggi, S. M., Buschkuhl, M., Jonides, J., and Gotlib, I. H. (2012). The everyday emotional experience of adults with major depressive disorder: Examining emotional instability, inertia, and reactivity. *Journal of Abnormal Psychology*, 121(4):819.
- Verbeke, G. and Lesaffre, E. (1996). A linear mixed-effects model with heterogeneity in the random-effects population. *Journal of the American Statistical Association*, 91(433):217–221.
- Wang, T., Merkle, E., and Zeileis, A. (2014). Score-based tests of measurement invariance: Use in practice. *Frontiers in Psychology*, 5(438):1–11.
- Wang, T. and Merkle, E. C. (2017). *merDeriv: Case-Wise and Cluster-Wise Derivatives for Mixed Effects Models*. R package version 0.1–1.
- Wang, T. and Merkle, E. C. (2018). Derivative computations and robust standard errors for linear mixed effects models in *lme4*. *Journal of Statistical Software*, 87(c01):1–16.
- Wang, T., Strobl, C., Zeileis, A., and Merkle, E. C. (2018). Score-based tests of differential item functioning via pairwise maximum likelihood estimation. *psychometrika*, 83(1):132–155.
- Zeileis, A. (2006). Implementing a class of structural change tests: An econometric computing approach. *Computational Statistics & Data Analysis*, 50(11):2987–3008.
- Zeileis, A. and Hornik, K. (2007). Generalized M-fluctuation tests for parameter instability. *Statistica Neerlandica*, 61:488–508.
- Zeileis, A., Leisch, F., Hornik, K., and Kleiber, C. (2002). *strucchange: An R package for testing structural change in linear regression models*. *Journal of Statistical Software*, 7(2):1–38.
- Zhang, X., Shao, X., Hayhoe, K., Wuebbles, D. J., et al. (2011). Testing the structural stability of temporally dependent functional observations and application to climate projections. *Electronic Journal of Statistics*, 5:1765–1796.