A Fine-Tooth Comb for Measurement Reliability: Predicting True Score and Error Variance in Hierarchical Models

Donald R. Williams, Stephen R. Martin,
Michaela C. DeBolt, Lisa M. Oakes, and Philippe Rast
University of California, Davis

The primary objective of this work is to extend classical test theory (CTT), in particular, for the case of repeated measurement studies. The guiding idea that motivates this work is that any theory ought to be expanded when it is not compatible with commonly observed phenomena–namely, that homogeneous variance components appear to be the exception and not the rule in psychological applications. Additionally, advancements in methodology should also be considered in light of theory expansion, when appropriate. We argue both goals can be accomplished by merging heterogeneous variance modeling with the central tenants of CTT. To this end, we introduce novel methodology that is based on the mixed-effects location scale model. This allows for fitting explanatory models to the true score (between-group) and error (within-group) variance. Two illustrative examples, that span from educational research to infant cognition, highlight such possibilities. The results revealed that there can be substantial individual differences in error variance, which necessarily implies the same for reliability, and that true score variance can be a function of covariates. We incorporate this variance heterogeneity into novel reliability indices that can be used to forecast group or person-specific reliability. These extend traditional formulations that assume the variance components are homogeneous. This powerful approach can be used to identify predictors of true score and error variance, which can then be used to refine measurement. The methods are implemented in the user-friendly R package **ICCier**.

Keywords: measurement reliability, Bayesian, classical test theory, heterogeneous variance, hierarchical model

It seems likely that in mental testing we shall find ourselves in the future increasingly concerned with measurement of fluctuating mental functions.

—(p. 340, Thouless, 1936)

A fundamental concept underlying psychological science is measurement reliability. From educational researchers to developmental psychologists, we all understand that gathering reliable measurements plays an important role in psychological inquiry in particular. This is because we often face the difficult task of investigating latent constructs that cannot be directly measured; for example, using a set of items to measure positive affect, rating mother-child interactions for the purpose of attachment classification, or investigating cognitive function experimentally, all share that common thread.

This challenge has resulted in a rich literature on psychological assessment, wherein measurement is often conceptualized with classical test theory (CTT, Thorndike, 1904; Traub, 1997).

In CTT, the observed measurements (X) are composed of a given individual's true score (T) plus a *random* error (E), resulting in $X = T + E$ (Lord & Novick, 1968a; Novick, 1965). With these ingredients, variability of observed measurements is defined as the sum of true-score variance, σ_T^2, plus random error variance, σ_E^2. Reliability is computed from these variances, that is,

$$r = \sigma_T^2 / \sigma_X^2 \tag{1}$$

which is by definition the proportion of the total observed variance, $\sigma_X^2 = \sigma_T^2 + \sigma_E^2$, due to the true score variance. Notice in Equation (1) that the true scores are the source of variation. Indeed, a critical assumption underlying CTT is that the errors (or residuals) are random and unsystematic, reflecting consistency over time of measurement, experimental condition, and individuals–homogeneous error variance (σ_E^2). This assumption may not be tenable in a wide range of psychological applications and thus randomness in the error structure could be the exception rather than the rule. As

Kline (2005) stated, "systematic errors are not handled well in CTT" (p. 94). We concur with Kline's assessment and argue that they are not handled at all.

Recently, the notion that heterogeneous error variance is omnipresent in repeated measurement studies has reached a crescendo, as systematicity of error has been found across psychological domains–it appears to apply generally. For example, Williams, Zimprich, and Rast (2019) examined the residual variance in a learning task, where it was found that the errors and their associated variance across repeated trials followed a nonlinear trajectory, including individual variation therein (see Figure 2 in Williams, Zimprich, & Rast, 2019). This reveals a problem with the basic assumptions of CTT, in that σ_E^2 is assumed constant across individuals and it cannot be a function of learning (or any other variable for that matter, Kline, 2005). Additionally, Williams, Liu, Martin, and Rast (2019) found that the residual variance of daily affect can be related to an individual's emotional state, as well as physical activity, on the previous day. This again suggests that the errors are not random.

These signatures in the error variance are not restricted to observational studies and have recently emerged in experimental settings. In particular, Williams, Martin, and Rast (2019) demonstrated that there was substantial individual variation in the error structure of cognitive inhibition tasks, which, they argued, translates into individual differences in the reliability of those tasks. That is, the participants were not accurately described by a common value for σ_E^2, which is assumed to be the case in Equation (1). These findings highlight an important problem in the way error is typically considered, because, in traditional reliability indices, each person (or more generally each group) is assumed to have a common error variance, which was clearly not the case in Williams, Martin, and Rast (see Figure 1 in 2019).

Once we acknowledge that error variance is not constant across measurement units but is heterogeneous, we can think of ways to *explain* it. The potential for modeling and explaining within-person variability has been acknowledged for some time and the current work is inspired by this body of research (i.e, the error, Cleveland, Denby, & Liu, 2003; Hedeker, Mermelstein, & Demirtas, 2008, 2012; Leckie, French, Charlton, & Browne, 2014; Rast & Ferrer, 2018; Rast, Hofer, & Sparks, 2012). A common finding is that the average within-person variance is not typically a good representation of any single individual's variance. For example, Williams, Martin, and Rast (2019, see Figure 3 therein) showed that the vast majority of individuals differed from the average within-person variance in a hierarchical (a.k.a. multilevel) model. The authors further noted that "...heterogeneous within-person variance is a defining feature of these [cognitive] tasks..." (p. 1). This is particularly relevant in the context of reliability because the *mere* average variability is being used to compute reliability (p. 35, Lord &

Novick, 1968b). Together, these findings suggest that error variance may be systematic but the implications for reliability are largely under-explored in psychological research.

We are not the first to propose a framework to overcome limitations in CTT. In particular, generalizability theory (*G*-theory) emerged to explicitly broaden how we think about and evaluate measurement reliability (Brennan, 1992; Cronbach, 1972; Shavelson & Webb, 2010). Similar to one of our primary goals, *G*-theory was motivated by the observation that "In classical test theory measurement error is undifferentiated random variation; the theory does not distinguish among various possible sources" (p. 599 Shavelson & Webb, 2012). This issue is addressed in *G*-theory by partitioning the variance into several components. However, *G*-theory still assumes that the error variance is "constant for all persons, regardless of true score" (p. 321, Shavelson & Webb, 2012), an observation that is inconsistent with the findings from (Williams, Martin, & Rast, 2019). To our knowledge, there is nothing intrinsic to either CTT or *G*-theory that necessitates assuming a common variance. This was noted in Gu, Emons, and Sijtsma (2018) "this [homogeneous error variance] is not an a priori assumption in CTT" (p. 4). Rather, homogeneity of variance is the standard assumption of commonly used statistical methodologies (e.g., ANOVA) used to compute reliability, including in both CTT and *G*-theory, out of convenience and then perhaps taken as a given. But, in general, as has been shown in work on within-person variability modeling, this assumption can and should be relaxed.

Consider the case of a one-way random effects model (Bartko, 1966), which is the focus of this current work. There are two sources of variation: the between- (σ_b^2) and the within-group (or person) (σ_w^2) variance component. From these components, we can compute the ratio

$$\rho = \frac{\sigma_b^2}{\sigma_b^2 + \sigma_w^2}. \tag{2}$$

This is commonly referred to as ICC(1), and it serves as a reliability index for single scores that ranges from $0-1$ (Shieh, 2016). In Equation (2), σ_b^2 is the variance in true scores or means. Reliability is then proportion of variance attributed to σ_b^2. Further, σ_w^2 is often referred to as measurement error. In cognitive inhibition tasks, for example, it captures the trial-to-trial "noise" in reaction times. Thus, assuming that σ_b^2 is held constant, increasing σ_w^2 will necessarily decrease reliability (Hedge, Powell, & Sumner, 2018). This definition of ICC, however, does not allow for the possibility of individual differences in reliability. Addressing this limitation was the primary focus of Williams, Martin, and Rast (2019), where it was demonstrated that there are individual differences in σ_w^2 and thus also in reliability. This insight is not possible with traditional approaches used to compute reliability. Building upon that finding forms the impetus for the current work.

We introduce novel methodology to probe measurement

reliability at both the level of the numerator and denominator in Equation (2). This is accomplished by extending Williams, Martin, and Rast (2019), where only a random intercepts model was fitted to σ_w^2, such that each person had their own error variance and person-specific reliability. In the current work, we now predict the "true" score (σ_b^2) and error (σ_w^2) variance with sub-models. The latter can also capture individual differences in error variance. This powerful approach for characterizing reliability allows for investigating whether σ_b^2 can be explained by other factors. For example, it may be true that older individuals are *more* homogeneous, which implies less true score variance and thus *less* reliable, assuming σ_w^2 is held constant. This general idea can be extended beyond variance between people. Suppose that σ_w^2 is permitted to vary among schools. Here, each school would then have their own error variance that can also be predicted by other variables such as gender or socioeconomic status. This would result in school-specific reliability that is also a function of covariates. These kinds of insights are possible by fitting explanatory models to the variance components in Equation (2).

The current work is organized as follows. In the first section, we introduce the proposed model. Our intention here is to describe key aspects of the proposed methodology. This serves as the foundation for the remainder of the paper. The rest of the work consists of case studies that span from educational research to cognitive psychology. This demonstrates the utility of the methodology. In this section, we also emphasize the connection between the presented methodology and classical test theory in hierarchical models. We end by summarizing our major and novel contributions, as well as discussing implications for measurement in psychological research.

Model Formulation

The presented methodology is based upon a straightforward extension of the traditional mixed-effects approach, which allows for partitioning the unexplained variance, or within-group variance, as well as the between-group variance, or true score variance. The technique to do so is termed mixed-effects *location scale* model (MELSM, pronounced mel·zəm, Hedeker et al., 2008, 2012), which combines earlier work on variance heterogeneity (Aitkin, 1987) and models for random scale effects (Cleveland et al., 2003). In the current work, we build upon this foundation to demonstrate that the MELSM has untapped potential as a fine-tooth comb for assessing measurement reliability.

Mixed-Effects Location Scale Model

The starting point is the standard linear mixed effects model for $i = 1, 2, \ldots, G$ groups (e.g., people or schools)

and $j (j = 1, 2, ..., n_i)$ measurements that may be specified as

$$\mathbf{y}_i = \mathbf{X}_i\boldsymbol{\beta} + \mathbf{Z}_i\mathbf{b}_i + \boldsymbol{\epsilon}_i, \tag{3}$$

where $\mathbf{y}_i$ is the $n_i \times 1$ response vector for observations in group i. $\mathbf{X}_i$ is the $n_i \times k$ design matrix for the fixed effects for observations in group i. Note that an intercept only results in the unconditional model (Raudenbush & Bryk, 2002), whereas including explanatory variables in $\mathbf{X}_i$ allows for computing conditional ICCs (Rabe-Hesketh & Skrondal, 2008). $\boldsymbol{\beta}$ captures the fixed effects and its dimension is $k \times 1$. The random effects are in the $n_i \times q$ matrix $\mathbf{Z}_i$ for observations in group i where $\mathbf{b}_i$ is the according $q \times 1$ vector with the random effects coefficients. These effects characterize the group means for the response (i.e., the location). $\boldsymbol{\epsilon}_i$ is a vector of errors specific to group i. The general assumption in standard mixed effects models is that random effects are $\mathbf{b}_i \sim N(\mathbf{0}, \boldsymbol{\Phi})$, where $\boldsymbol{\Phi}$ is a $q \times q$ covariance matrix for the random effects with the variances σ_b^2 and the covariances $\sigma_{bb'}$ (for $q \neq q'$). The errors $\boldsymbol{\epsilon}_i$ are also assumed to be normally distributed with a mean of $\mathbf{0}$ and covariance of $\sigma_\epsilon^2\boldsymbol{\Psi}_i$ where $\boldsymbol{\Psi}_i$ is a $n_i \times n_i$ matrix which can take different structures. In these models the between-group variance is captured by σ_b^2 and the within-group variance is represented in σ_w^2. In the context of reliability, the former is termed true score variance. This formulation leads to computing traditional measures of reliability, for example, ICC(1) given in (2)–that is, $\sigma_b^2/\sigma_b^2+\sigma_w^2$.

Error Variance. In this standard form, the error variance σ_ϵ^2 is a fixed entity. To allow it to differ at the group level, we add the subscript i to the within-group variance term (cf. Hoffman, 2007; Myles, Price, Hunter, Day, & Duffy, 2003) and we also allow it to differ among j-measurements to obtain $\sigma_{\epsilon_{ij}}^2$. Changes in the within-group variance $\sigma_{\epsilon_{ij}}^2$ are explained by group-varying covariates in the $n_i \times m$ matrix $\mathbf{W}_i$ for the fixed effects and $\mathbf{V}_i$, with dimension $n_i \times p$ (and $m \geq p$) for the random effects (Rast et al., 2012). Hence, with the inclusion of level one covariates, the within-group variance not only varies across groups but also across measurements given the model:

$$\boldsymbol{\varphi}_i = \exp(\mathbf{W}_i\boldsymbol{\eta} + \mathbf{V}_i\mathbf{t}_i). \tag{4}$$

$\boldsymbol{\varphi}_i$ then is the $n_i \times 1$ vector that contains all error variances $\sigma_{\epsilon_{ij}}^2$ for group i and for each measurement j within that group. $\boldsymbol{\eta}$ is comparable to the regression weights $\boldsymbol{\beta}$ in (3). That is, for an intercept and slope term, η_0 defines the average within-group variance and η_1 weights the influence of the predictor on the variance. The individual departures from the fixed effects that are captured in the random effects $\mathbf{t}_i$ are normally distributed with $\mathbf{t}_i \sim N(\mathbf{0}, \boldsymbol{\Theta})$, where $\boldsymbol{\Theta}$ is a covariance matrix of dimension $p \times p$ that contains the random effects of the scale. Note that $\mathbf{W}_i$ and $\mathbf{V}_i$ may, or may not be the same as $\mathbf{X}_i$ and $\mathbf{Z}_i$. The exponent is used to ensure that the variance is restricted to positive values, and thus, is log-normally

distributed (Hedeker et al., 2008). Together, this effectively allows the denominator in Equation (2) to be a function of it own mixed-effects sub-model.

True Score Variance.   A key contribution of this work is predicting true score variance. Accordingly, we also introduce a sub-model for the between-group variance. This permits the true score variance in Equation (2) to be a function of its own sub-model. It is important to note here that there are now random effects $\mathbf{b}_i$ from the *location* of the model (the means structure) and random effects $\mathbf{t}_i$ from the *scale* of the model (the within-group variance structure). All these random effects are assumed to come from a normal distribution with mean zero. Hence, we can stack both $\mathbf{b}_i$ and $\mathbf{t}_i$ vectors, resulting in $\mathbf{u}_i \sim N(\mathbf{0}, \Sigma_i)$. This also means that Σ_i contains the variances and covariances of both the location and scale. To define a variance model for Σ_i, we can decompose $\Sigma_i = \boldsymbol{\tau}_i \Omega \boldsymbol{\tau}'_i$, where $\boldsymbol{\tau}_i$ is a diagonal matrix for group i in which the diagonal elements are the random-effect standard deviations (that includes σ_b^2 in Equation 2) and Ω is the correlation matrix that contains the correlations among *all* random effects. That is, Ω is of dimension $(q+p) \times (q+p)$ and contains the correlations among the location, scale, and location and scale random effects. We can now define the sub-model for the random effects SDs, that is,

$$\mathrm{diag}(\boldsymbol{\tau}_i)' = \exp(\boldsymbol{g}_i \boldsymbol{\iota}), \qquad (5)$$

where $\boldsymbol{g}_i$ is the design matrix that contains between-group predictors (e.g., age group or gender), and $\boldsymbol{\iota}$ is a matrix with $(q+p)$ columns of coefficients. ι_{rc} is the effect of the rth column in $\boldsymbol{g}_i$ on the cth random effect SD (i.e., $\tau_{c,c}$). For example, ι_{12} is the intercept for $\tau_{2,2}$, and ι_{23} is a slope parameter for $\tau_{3,3}$. Note that our focus is a one-way model with random intercepts for both the location or means ($\tau_{1,1}$) and the scale or within-group variance ($\tau_{2,2}$). Consequently, the random-effects variance is not constant but may change due to group-specific characteristics (e.g. Leckie et al., 2014). If there is a positive effect of age, this would mean that older people are relatively more heterogeneous in their true scores and thus also more reliable.

Having specified all elements, we can define the full MELSM as

$$y_i \sim N(\mu_i, \boldsymbol{\varphi}_i)$$
$$\mu_i = \mathbf{X}_i \boldsymbol{\beta} + \mathbf{Z}_i \mathbf{b}_i$$
$$\boldsymbol{\varphi}_i = \exp(\mathbf{W}_i \boldsymbol{\eta} + \mathbf{V}_i \mathbf{t}_i)$$

with the random effects for both the location and the scale coming from the same multivariate distribution

$$\begin{bmatrix} \mathbf{b}_i \\ \mathbf{t}_i \end{bmatrix} \sim N(\mathbf{0}, \Sigma_i)$$
$$\Sigma_i = \boldsymbol{\tau}_i \Omega \boldsymbol{\tau}'_i$$
$$\mathrm{diag}(\boldsymbol{\tau}_i) = \exp(\boldsymbol{g}_i \boldsymbol{\iota}).$$

The remaining "issue" is the prior distributions. In all models, we assign each parameter a diffuse prior, such that they are expected to have minimal impact on the estimates.

Group-Specific Reliability

Modeling the variance structure leads to group-specific reliability. These groups can be people, schools, or countries, that is, some higher level in which observations are nested within. First consider the case of a random intercept models fitted to both the location and scale, which provides the population-average within-group variance (i.e., the fixed effect) and random effects for each group. In other words, each group has their own error variance. Group-specific reliability can the be computed with a straightforward extension to ICC(1) given in Equation (2), that is,

$$\rho_i = \frac{\tau_{1,1}^2}{\tau_{1,1}^2 + \exp[\eta_0 + t_i]^2}. \qquad (6)$$

Note that the subscript i denotes the ith group, $\tau_{1,1}^2$ is the true score variance (i.e., σ_b^2), and $\exp[\eta_0 + t_i]^2$ is the group-specific variance. More specifically, with $i = 1$, this formulation would provide the group-specific estimate of reliability for the first group (or person). Furthermore, in Equation (6), the covariance between two observations from the same group remains unchanged from the customary definition of ICC(1). In other words, the modification is slight, in that the correlation is expressed as a function of the within-group variances (in standard mixed-effects models there is only one within-group variance; that is, the average in our model).

Predicting Reliability

The above formulation is restricted to estimating group-specific reliability. The central idea behind the current work is to also fit explanatory models to both the true score (σ_b^2) and error (σ_w^2). This allows for predicting reliability and it extends customary theories of measurement, where the variance components are taken to be homogeneous. To this end, we follow Hedeker et al. (2008) and define the covariance for any two measurements from the same group as

$$\mathrm{Cov}(y_{ij}, y_{ij'}) = \sigma_{bi}^2 \qquad (7)$$
$$= \tau_{i1,1}^2 = \exp(\boldsymbol{g}_i \boldsymbol{\iota}_1)^2 \text{ for } j \neq j'. \qquad (8)$$

Note that the subscript i is used to emphasize that the variance is a function of the design matrix $\boldsymbol{g}_i$ and the regression weight vector $\boldsymbol{\iota}_1$. For each group, the predicted reliability then takes on the following form

$$\rho_i = \frac{\exp(\boldsymbol{g}_i \boldsymbol{\iota}_1)^2}{\exp(\boldsymbol{g}_i \boldsymbol{\iota}_1)^2 + \exp(\mathbf{W}_i \boldsymbol{\eta} + \mathbf{V}_i \mathbf{t}_i)^2} \qquad (9)$$

Hence, reliability can be probed at both the level of the numerator and the denominator of Equation (2). Of course, if there is not much group variability in the variance structure, and if the covariates have a minimal effect on the variances, this would result in Equation (2) and Equation (9) producing similar estimates. This is because a mixed-effects model is a special case of the MELSM, but with an implicit fixed intercept only model fitted to the variance σ_w^2 and homogeneous random effect variances.

Note that the above is meant to be generic, which highlights the generality of our approach for assessing measurement reliability. In the case studies, we provide specific models and reliability indices that will further clarify this formulation.

Connection to the Broader Literature

Before proceeding to the illustrative examples, we want to emphasize the connection of this work to the broader literature. Although varying ICCs naturally arise from heterogeneous variance modeling, these ideas are most prominent in research areas that gather intensive longitudinal data (Hamaker, Asparouhov, Brose, Schmiedek, & Muthén, 2018; Hedeker et al., 2012; Rast & Ferrer, 2018; Watts, Walters, Hoffman, & Templin, 2016; Williams, Liu, et al., 2019). Indeed, to our knowledge, the notion was first described in the context of ecological momentary assessment. In particular, Hedeker et al. (2008) described how the variances (e.g., σ_b^2 and σ_w^2) could be a function of covariates, providing the foundation for work by Brunton-Smith, Sturgis, and Leckie (2017) and Williams, Martin, and Rast (2019). These works in particular estimated group-specific ICCs in hierarchical models.

There are several novel aspects of the present work. First, we *fully* merge the ideas stemming from the variance modeling literature with psychological measurement. In particular, we provide the key insight that the MELSM is ideal for assessing and *predicting* reliability. The implications of this are far reaching. First, the long-standing issue of assuming a constant within-group variance is addressed. Second, we adopt a Bayesian framework for extending CTT. The advantages of Bayesian methods for estimating true score and error variance were described in Lindley (1969) and Novick, Jackson, and Thayer (1971). Our work builds upon those ideas, but with the full power of modern Markov chain Monte Carlo algorithms (Betancourt, 2017), which were not computationally feasible until the 1990's (e.g., see for example Gelfand & Smith, 1990; Robert & Casella, 2011). Third, we demonstrate the utility of this framework in several psychological applications. Lastly, we have implemented this methodology in the user-friendly R package **ICCier**, which serves as a high-level interface to the programming language Stan (Stan Development Team, 2016).

Illustrative Examples

In this section, we employ the methodology in two cases that represent a wide range of psychological applications. Our intention here is to demonstrate the utility of heterogeneous modeling for assessing reliability. Recall that the motivation for this work was partly based on findings demonstrating that homogeneity in the variance structure appeared to be the exception rather than the norm in psychological applications. Thus, in these two examples, we start by characterizing the within-group variance and then proceed to predicting reliability. This necessarily requires formulating a variety of models. However, it is important to note that the central ideas from each case study apply generally.

Case 1: Educational Research

Educational researchers often encounter hierarchically structured data. That is, while students serve as the unit of measurement, they are commonly nested within schools that can in turn be nested within a higher-level unit, such as the respective school district, county, or even at the state level. It is for this reason educational data is often used to demonstrate key aspects of hierarchically modeling. To our knowledge, however, the MELSM has only been used once to investigate within-school variance heterogeneity (Leckie et al., 2014).

An important question in education is to identify factors that are related to student success and academic achievement. Such factors include school location (Logan & Burdick-Will, 2017), teacher expectations (de Boer, Timmermans, & van der Werf, 2018), and the socioeconomic status composition among attending students (Sirin, 2005). However, reducing education inequality between, for example, suburban and urban schools (Sandy & Duncan, 2010), has proven to be far from trivial in that " ...income-related gaps both in access to and in success in higher education are large and growing" (p. 125, Haveman & Smeeding, 2006). Consequently, an important question is to not only study differences between developed environments, but to also investigate whether certain schools differ from each other in test scores. This can provide important information at the school-level (e.g., rankings) and it is the focus of the following example.

Average Score Reliability. Assessing school-level differences in average test scores is inextricably linked to average score reliability. In hierarchical models, an estimate of reliability guides the degree to which the estimator relies on observed group means or the average across groups (Raudenbush & Bryk, 2002). That is, in the presence of large error relative to between-group variance, the means are smoothed towards the grand mean (the groups become more similar to one another). This is sometimes referred to as partial pooling (e.g., Gelman, Hill, & Yajima, 2012). The degree of smooth-

ing for each group i is defined as

$$\lambda_i = \frac{\sigma_b^2}{\sigma_b^2 + \sigma_w^2/n_i}, \tag{10}$$

which is known as the "shrinkage factor" and it is *equivalent* to average score reliability or ICC(2) (McGraw & Wong, 1996; Shieh, 2016). In Equation (10), n_i is the number of measurements for group i. Note that this is the same as ICC(1) given in Equation (2), but, in this case, σ_w^2 is divided by the respective number of measurements gathered from a given group. This establishes the foundational link to CTT, in that reliability approaches one as $n_i \rightarrow \infty$, with the assumption that σ_b^2 is greater than zero. The predicted estimate $\hat{\mu}_i$, or the average test score, for school i is then computed with

$$\hat{\mu}_i = (1 - \lambda_i)\bar{Y}_G + \lambda_i\bar{Y}_i, \tag{11}$$

where $\bar{Y}_G$ is the grand mean and $\bar{Y}_i$ is the school-specific mean (McCulloch, 2003). This is commonly referred to as the Best Linear Unbiased Predictor (BLUP) and it is related to the so-called "Stein's paradox" (Efron & Morris, 1977; Morris & Lysy, 2012; Stein, 1956). However, it is important to note that the original derivation is grounded in psychometrics and CTT. That is, although Equation (11) was given as an individuals best estimate of their true score in Kelley (p. 178, see eq. 22, 1927), it makes the critical assumption of a common variance (i.e., σ_w^2 in Equation 10).

Herein lies how our methodology can be viewed as an extension to CTT. Namely, in Equation (10), the calculation of average score reliability assumes a common within-group variance σ_w^2. This is not so for the MELSM, and in reference to group-specific reliability given in Equation (6), we can further establish a connection to the Spearman-Brown prediction formula (Brown, 1910; Spearman, 1910), that is,

$$\Psi(\rho)_i = \frac{J \cdot \rho_i}{1 + (J - 1)\rho_i}. \tag{12}$$

Here the Spearman-Brown equation is applied to the group-specific reliability, ρ_i, based on ICC(1), such that each school has their own predicted reliability for J measurements (the group size). Typically the Spearman-Brown formulation assumes a common value for ρ, that is, the subscript i does not exist. Note that Equation (12) approximates ICC(2), that was given in Equation (10), asymptotically as J increases (Bliese, 2000). Because the "shrinkage factor" is equivalent to ICC(2), the Spearman-Brown equation in Equation (12) can also be understood as predicting the expected degree of shrinkage in a one-way random effects model. The generalization to predicting group-specific reliability and the connection to shrinkage are major contributions of this work.

Model Specification. We now apply the MELSM and the Spearman-Brown equation for predicting reliability

based on J measurements in an educational setting. The basic idea is to predict reliability, in an effort to reach an acceptable level while also accounting for (possible) heterogeneity in the within-school variance structure. Further, assuming the goal is to detect school-level differences, Equation (12) can also be used to predict the expected shrinkage. This is critical to detect differences, in that, with increasing shrinkage, the school-level means or true scores become more homogeneous (Gelman et al., 2012).

We use data from the General Certificate of Secondary Education (GCSE) exam, which is an academic qualification for the United Kingdom. There are 65 schools and 4,059 students from six Inner London Education Authorities. These data have been used in several examples demonstrating the utility of hierarchical models. In particular, the school-level residuals have also been examined in an attempt to make inference about the variance structure (Goldstein et al., 1993, see Figure 2,), which is a naïve MELSM. In the present analysis, the outcome is normalized test scores at age 16. We also predict true score variance, σ_b^2, with school-level averages for the London Reading Test (LRT) at age 11. The schools have been categorized into three groups based on the average student level Verbal Reasoning (VR) score at intake, consisting of the bottom 25%, middle 50%, and top 25%.

Mean Structure. For the ith school and jth measurement, the one-way random effects model is defined as

$$y_{ij} = \beta_0 + u_{0i} + \epsilon_{ij}, \tag{13}$$

where β_0 is the fixed effect and u_{0i} the individual deviation. More specifically, β_0 is the average of the school means or observed scores and $\beta_0 + u_{01}$ is the mean for the ith school.

"True" Score Variance. The random effects are then assumed to be drawn from a Gaussian distribution, that is,

$$u_{0i} \sim \mathcal{N}(0, \sigma_{bi}^2).$$

Here the between-school variance σ_b^2 captures the variability in the random effects $var(u_{0i})$, or the true score variance, that are assumed to be normally distributed with a mean of zero. We then predict the true score variance with the school-level rankings of test scores. This log-linear model is defined as

$$\sigma_{bi}^2 = \exp(\kappa_0^{(-25\%)} + \kappa_1^{(50\%)}X_{i1} + \kappa_2^{(+25\%)}X_{i2}). \tag{14}$$

X is the design matrix with two dummies, coded to represents the schools. Thus $\kappa_0^{(-25\%)}$ is the between-school variance for the schools in the bottom 25th percentile in scores at age 11. This serves as the reference category or the intercept. Consequently, for the schools in the middle 50th and top 25th percentile, the variance in their true scores is $\kappa_0^{(-25\%)} + \kappa_2^{(+50\%)}$ and $\kappa_0^{(-25\%)} + \kappa_2^{(+25\%)}$, respectively.

Error Variance. Further, given the MELSM approach, the residuals or errors are also assumed to be normally distributed with a mean of zero and variance $\sigma_{\epsilon_{ij}}^2$, given as a function of a linear model, such that

$$\epsilon_{ij} \sim \mathcal{N}(0, \sigma^2_{\epsilon_{ij}}) \text{ with}$$
$$\sigma^2_{\epsilon_{ij}} = \exp(\eta_0 + u_{1i})^2.$$

As indicated by the subscripts i and j, the error variance $\sigma^2_{\epsilon_{ij}}$, is now allowed to vary across i schools and j measurements given a log-linear model. The parameters in the scale model (the model for the error variance) are analogous to those in Equation (16). η_0 represents the intercept and defines the average of the within-school variances and u_{1i} represent the random effect, that is, the school departures from η_0. For the first school ($i = 1$), $\eta_0 + u_{11}$ is the variability of their exam scores. The random school effects are also normally distributed, $u_{1i} \sim \mathcal{N}(0, \sigma^2_1)$, and like the mean structure, they can be shrunken towards the average within-school variance η_0.

A simple example using the **ICCier** package is available in the Appendix B.

Reliability Indices. Together, this model allows for computing group-specific reliability, given the within-school (error) and between-school (true score) variance. That is, for school i their respective reliability, or ICC(1), is

$$\rho_i = \begin{cases} \dfrac{\exp(\kappa_0)^2}{\exp(\kappa_0)^2 + \exp(\eta_0 + u_{1i})^2} & \text{if } -25\% \\[2ex] \dfrac{\exp(\kappa_0 + \kappa_1)^2}{\exp(\kappa_0 + \kappa_1)^2 + \exp(\eta_0 + u_{1i})^2} & \text{if } 50\% \\[2ex] \dfrac{\exp(\kappa_0 + \kappa_2)^2}{\exp(\kappa_0 + \kappa_2)^2 + \exp(\eta_0 + u_{1i})^2} & \text{if } +25\% \end{cases} \quad (15)$$

This can then be used to predict the average score reliability with Equation (12), which as we explained above, also predicts the expected shrinkage in a one-way random effects model. This is accomplished by plugging ρ_i from Equation (15) directly into Equation (12)

Results. We first compared the MELSM to a traditional, location only, random intercepts model. This is the traditional "CTT model." We used an approximate measure of predictive accuracy, that is, leave-one-out cross-validation (Vehtari, Gelman, & Gabry, 2017). The difference was nearly 6 standard errors away from zero, which indicates that the MELSM was preferred. This can interpreted as "rejecting" the CTT model for computing reliability, given there is heterogeneous variance.

Error Variance. This improved model fit with the MELSM can also be inferred from Figure 1. Panel A includes the within-school standard deviations, and the black bars are 90% credible intervals that excluded zero (the average). In total, 26% of the schools differed from the average within-school variance. This is important. Recall that this violates the implicit assumption of traditional reliability indices, such as for ICC(1) defined in Equation (2). Further, panel A also highlights a key aspect of the model. Namely,

there is shrinkage towards the average within-school variance which, in some sense of the word, provides a more "reliable" estimate of the error. Consequently, the group-specific estimates of reliability defined in Equation (15) are computed from the partially pooled estimates of the school-specific variances.

"True" Score Variance. Panel B includes the estimated between-school variances. This is another key aspect of the proposed methodology, in that each category of school-level test scores has their own estimate of true score variance. Here it can be seen that the middle group, that is those in the middle 50% of LRT scores, were more homogeneous than the top 25% group. For a point of reference, the 90% credible interval (CrI) for $\kappa_1^{(50\%)}$ excluded zero, thereby indicating a difference. This necessarily translates into the measure being less reliable for middle 50% group. The implications of this cannot be understated, in that, when assuming there are homogeneous variance components, this can mask important information related to achieving a desired level of reliability at the school level.

Reliability. Panel C includes the school-specific reliability estimates, or ICC(1), computed with Equation (15). Recall that ICC(1) captures the correlation between any two observations from the same school. Importantly, these estimates are a combination of Panel A, that includes the within-school variance, and Panel B, that includes the between-school variance. They have been separated according to their test score ranking and the dotted line denotes the fixed effect or average reliability. The results are striking. For all schools in the middle 50% group, their respective 90% CrIs excluded the average reliability in the other categories. Indeed, their average reliability was merely 0.10, whereas it was 0.25 and 0.30 in the bottom 25% and top 25% groups, respectively. Said another way, the percentile rank of test scores *explains* reliability in these schools. This insight was made possible by fitting sub-models to both the between (true score) and within-school (error) variance structures.

We next investigated average score reliability, predicted with Equation (12), while accounting for heterogeneity in the variance structures. The idea is to determine the number of measurements needed to achieve adequate reliability. Figure 1 (panel D) displays these results. Note that each line represents a school. As a point of reference, we also included a location only model that embodies the assumptions of CTT. The critical difference is that the variance components in Equation (2) are fixed or homogeneous. For the CTT model, the black line is the predicted reliability and the grey ribbon is the corresponding 90% CrI. The results are again striking. Namely, for the middle 50% category, most of the schools were far below the predicted reliability of the CTT model. For example, on the one hand, the CTT model would have us believe that perhaps less than 25 measurements are needed to obtain, say, an average score reliability of 0.80. On the other
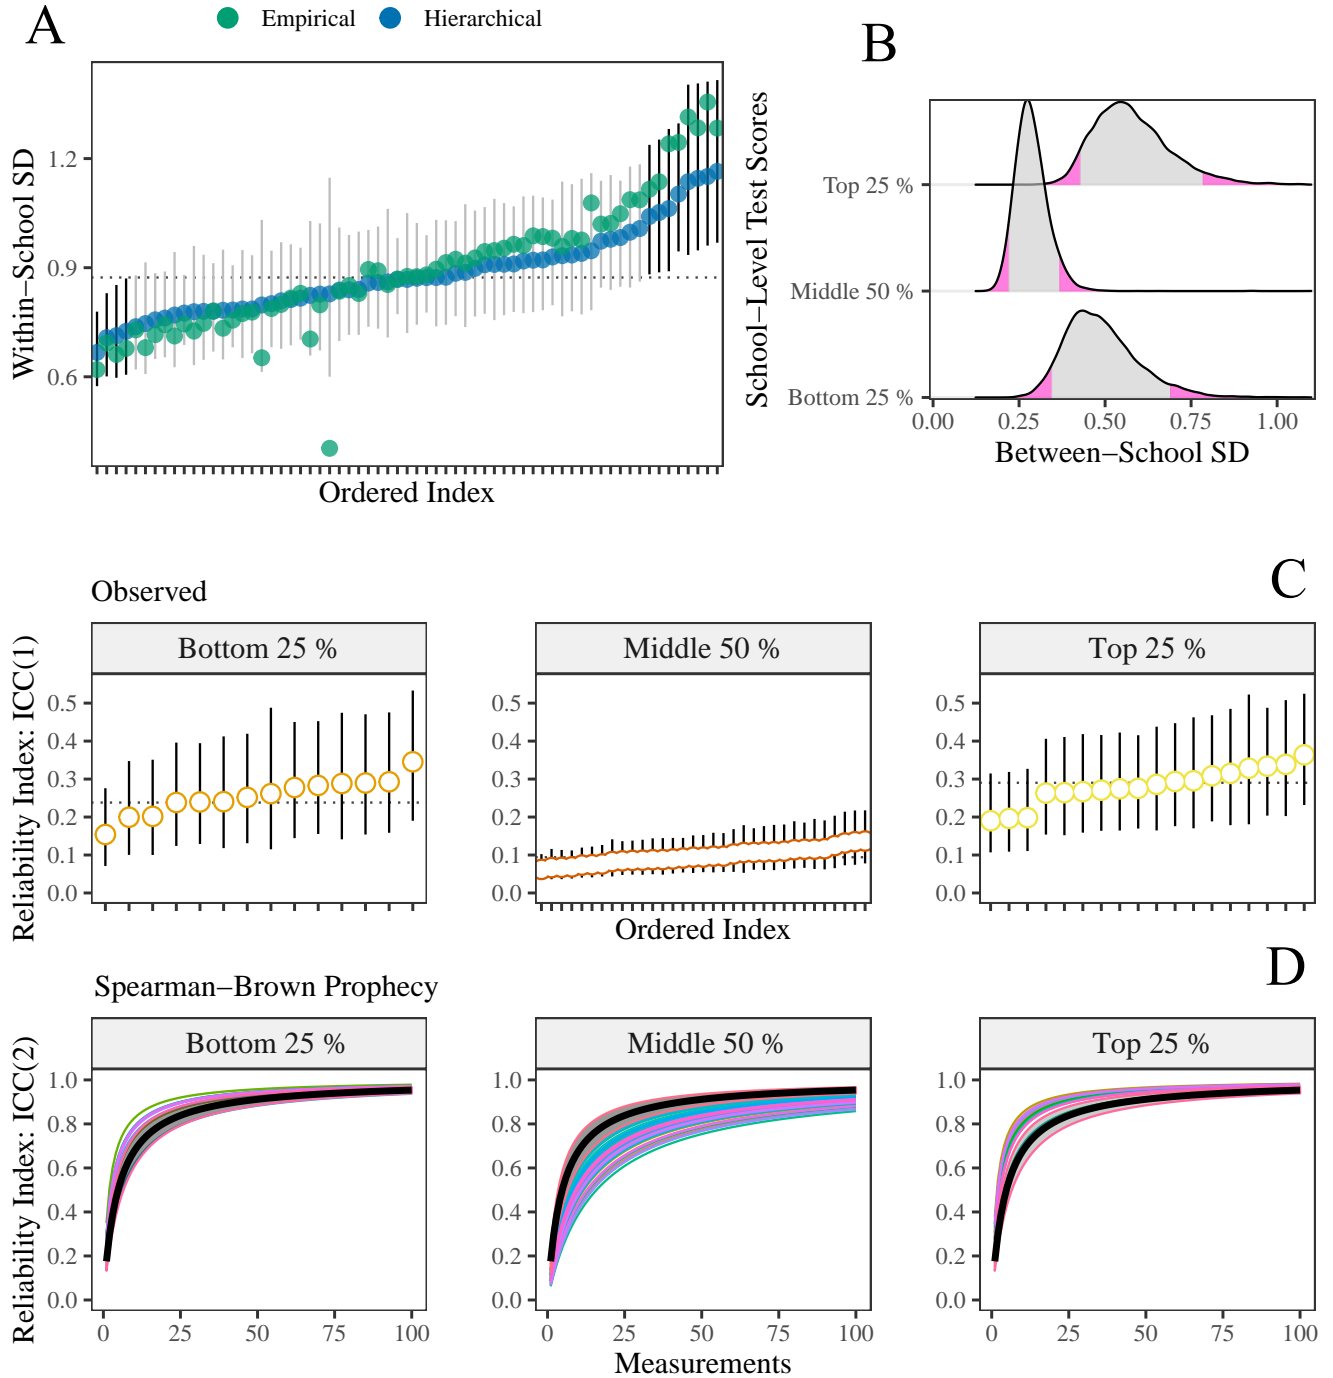
Figure 1. Results for Case 1: Educational Research. A) Within-school standard deviations (*SD*). The error bars are 90% credible intervals (CrI). The black error bars excluded the dotted line that denotes the average, or fixed effect, within-school *SD*. The blue points are the shrunken estimates that have been smoothed towards the fixed effect. The green points are the empirical, non-shrunken, within-school *SD*s. B) Posterior densities for between-school (true score) variability according to school-level rankings. The grey area is the 90% highest density interval. C) School-specific reliability that are separated according to school-level rankings. ICC(1) is for single observations.The dotted lines denote the average reliability for each group. D) School-specific reliability predicted with the Spearman-Brown equation (12). ICC(2) is for mean scores. Each school is represented by a colored line. The black line is the classic "CTT model" that assumes the variance components (i.e., true score and error variance) are homogenous. The grey region is the 90% CrI for the "CTT model."

hand, by fitting sub-models to the variance components, it is clear that we might actually need 100 measurements. This is attributed to the smaller true score variance for schools that placed into the middle 50% of test scores (Figure 1, panel B).

Moreover, the results in panel D also translate into minimizing shrinkage towards the grand mean (Equation 11). This is important to consider when investigating between school mean differences. This is a result of the equivalence between average score reliability, or ICC(2), and the "shrinkage factor" given in Equation (10). Consequently, these results can also be used to inform study design: Based on these results, it is clear that more measurements would be needed from the middle 25% percentile group.[1] More generally, this points towards considering that the reliability of a test can vary at both the level of true score and error variance, that is, the numerator and denominator in Equation (2) for ICC(1).

Case 2: Developmental Psychology

As a second case we consider an example from infant cognitive development. Developmental researchers have long been interested in measuring individual differences with the desire to predict future behavior or developmental outcomes (Colombo & Fagen, 2014). For example Rose, Feldman, and Wallace (Rose, Feldman, & Wallace, 1988) found that performance on visual recognition tasks in a sample of 6-, 7-, and 8-month-old infants predicted later Stanford-Binet IQ scores at 3 years of age.

A prerequisite of psychological tasks used to measure individual differences is their ability to yield reliable measurements from participants. That is, repeated measurements obtained from a single individual must be correlated (to some extent) with other measurements obtained from the *same* individual on the same task. Recent work has drawn attention to the observation that several commonly used cognitive inhibition tasks fail to produce reliable measurements in adults (Hedge et al., 2018; Rouder, Kumar, & Haaf, 2019). As a consequence, the ability for these tasks to measure individual differences has been called into question.

The finding that commonly used cognitive inhibition tasks are not well suited to study individual differences is surprising and has generated interest in exploring whether these findings also apply to tasks used in developmental psychology (Arnon, 2018). However, it should be noted that conclusions that tasks do not reveal individual differences are drawn from models that assume a common variance between individuals. As previously demonstrated, the assumption of a common variance does not hold in many observed adult data sets (Williams, Martin, & Rast, 2019). Hence, a core assumption of measurement reliability is routinely violated (i.e., homogeneous variance), and to date a solution to address the so-called "reliability paradox" has not been proposed. Our methodology not only accommodates heterogeneous variance, but as we show below, opens the door for

explaining aspects of reliability–that is, true score and error variance. This can provide novel insights into cognitive tasks that could then be used to improve measurement.

Model Specification. We now apply the MELSM to characterize performance on a task of visual attention in infants (described in more detail below). There is emerging evidence that heterogeneous within-person variance is a defining feature of cognitive tasks. However, to date, this has only been investigated in college-aged adults. Hence, in this case study, we first characterize within-infant variability and also infant-specific reliability in an experimental task that has four conditions. We then dive deeper into the reliability of those conditions. Here we first look at how the number of trials completed influences reliability and then we assess reliability differences between conditions.

The data are drawn from an assessment of 5- to 12-month-old infants using an adaptation of the Infant Orienting With Attention (IOWA) task developed by Ross-Sheehy and colleagues (Ross-Sheehy, Schneegans, & Spencer, 2015). Eye gaze data were collected (using an automatic infra red light eye tracking system) from 98 full term typically developing infants in an attention cueing task. Each trial begins with a central fixation stimulus (a looming smiley face paired with classical music). As infants fixate that stimulus, a 100 ms spatial attention cue (a small black dot) is presented left or right of midline, followed by a target (a realistic photograph of an object, e.g., rattle or banana). In the task, infants received four types of trials: valid cue trials in which the target appeared in the location of the cue, invalid cue trials in which the target appeared on the opposite side of the cue, double cue trials in which two cues appeared followed by a single target, and no cue baseline trials in which no cue appears before the target appeared. Infants provided data on 50 trials on average, approximately 11 trials per condition. This is more data than is typically collected per infant in most paradigms, allowing us the opportunity to address the questions of interest in this work.

Mean Structure. Spatial attention and orienting speed were assessed using reaction time, or latency to fixate the target on each trial. We fitted separate models for each of the four conditions. For the ith infant and jth measurement, the one-way random effects model is defined as

$$y_{ij}^{(c)} = \beta_0^{(c)} + u_{0i}^{(c)} + \epsilon_{ij}^{(c)}. \tag{16}$$

Note that $(c) = \{1, 2, 3, 4\}$ is used to denote the respective experimental condition (i.e., a separate model was fitted for each). Hence $\beta_0^{(c)}$ is the fixed effect and $u_{0i}^{(c)}$ the individual deviation for a given condition c. More specifically, $\beta_0^{(c)}$ is the

---

[1]In these data, there was a large correlation ($r \approx 0.70$) between ICC(2) and the difference between the empirical and shrunken mean estimates. Hence, the less reliable schools tended to also be shrunken more so towards the grand mean.

average of the infants means or observed scores and $\beta_0^{(c)} + u_{01}^{(c)}$ is the mean for the ith infant.

"True" Score Variance. The random effects are then assumed to be drawn from a Gaussian, that is,

$$u_{0i}^{(c)} \sim \mathcal{N}(0, \sigma_{bi}^{2(c)}).$$

Here the between-infant variance $\sigma_{bi}^{2(c)}$ captures the variability in the random effects $var(u_{0i})^{(c)}$, or the true score variance, that are assumed to be normally distributed with a mean of zero. We then predict the true score variance with the number of trials completed by each infant. This explores the possibility that reliability is a function of trials completed. This model for the true score variance is defined on the log-scale as

$$\sigma_{bi}^{2(c)} = \exp(\kappa_0^{(c)} + \kappa_1^{(c)} X_i). \tag{17}$$

The predictor X, that is number of trials per infant, was grand mean centered. As a result, $\kappa_0^{(c)}$ represents the true score variance for the average number of trials across all infants. We also scaled the predictor such that $\kappa_1^{(c)}$ corresponds to the regression weight for a 10 trial increase. Now, rather than assuming σ_b^2 to be fixed and non-varying across all infants, as in CTT, our formulation allows for explaining true score variance in a cognitive task–in this case with the number of trials completed.

Error Variance. Further, with the MELSM approach, the residuals or errors are also assumed to be normally distributed with a mean of zero and variance $\sigma_{\epsilon_{ij}}^2$, given as a function of a linear model, such that

$$\epsilon_{ij}^{(c)} \sim \mathcal{N}(0, \sigma_{\epsilon_{ij}}^{2(c)}) \text{ with} \tag{18}$$

$$\sigma_{\epsilon_{ij}}^{2(c)} = \exp(\eta_0^{(c)} + u_{1i}^{(c)})^2. \tag{19}$$

As indicated by the subscripts i and j, the error variance $\sigma_{\epsilon_{ij}}^{2(c)}$, is now allowed to vary across i infants and j measurements given a log-linear model. The parameters in the scale model (the model for the error variance) are analogous to those in Equation (16). That is, $\eta_0^{(c)}$ represents the intercept and defines the average of the within-infant variances and $u_{1i}^{(c)}$ represent the random effect, that is, the individual departures from $\eta_0^{(c)}$. For infant $i = 1$, $\eta_0^{(c)} + u_{11}^{(c)}$ is the variability of their latencies in condition c. The random effects are also normally distributed, $u_{1i}^{(c)} \sim \mathcal{N}(0, \sigma_1^{2(c)})$, and like the mean structure, they can be shrunken towards the average within-infant variance $\eta_0^{(c)}$.

Reliability Indices. With the model specification in hand, we can now define reliability indices for each condition, c, and infant, i. This is defined as

$$\rho_i^{(c)} = \frac{\exp(\kappa_0^{(c)} + \kappa_1^{(c)} X_i)^2}{\exp(\kappa_0^{(c)} + \kappa_1^{(c)} X_i)^2 + \exp(\eta_0^{(c)} + u_{1i}^{(c)})^2}. \tag{20}$$

This index is infant-specific ICC(1) and X_i corresponds to a given number of trials completed for infant i. The true score variance, or the numerator in Equation (20), is then a function of trials completed. Consequently, $\rho_i^{(c)}$ is computed with respect to the number of trials completed for each infant and their respective error variance, $\exp(\eta_0^{(c)} + u_{1i}^{(c)})^2$. Note that trials completed is not an individually varying predictor, but is instead a so-called level two variable. This allows for predicting reliability, given the fixed effect average for the within-infant variance. This predicted reliability index is defined as

$$\hat{\rho}^{(c)} = \frac{\exp(\kappa_0^{(c)} + \kappa_1^{(c)} \cdot \text{Trials})^2}{\exp(\kappa_0^{(c)} + \kappa_1^{(c)} \cdot \text{Trials})^2 + \exp(\eta_0^{(c)})^2}. \tag{21}$$

Recall from Equation (17) that $\kappa_0^{(c)}$ is the between-infant variance for the average number of trials completed (i.e., the grand mean) and $\kappa_1^{(c)}$ the corresponding regression weight. This allows for predicting reliability across a range of trials completed which translates into investigating the effect of completing more trials on reliability at the level of true score variance. Further, because the package **ICCier** employs Bayesian estimation, $\hat{\rho}^{(c)}$ has a full distribution which readily allows for assessing reliability differences, for example, between conditions and as a function of a level two predictor. This is demonstrated below. This powerful approach for probing reliability is a major and novel contribution to the measurement literature.

Results. We first compared the MELSM to a traditional, location only, random intercepts model. This was again done with LOO (Vehtari et al., 2017). The differences ranged from at most 9.4 (control condition) to 6.1 standard errors (double condition) away from zero, which indicates that the MELSMs had superior model fit compared to the customary "CTT model." This also implies that there were considerable individual differences in the true scores and within-person variability.

Error Variance. The improved fit of the MELSM can be inferred from Figure 2 (panel A), which includes the within-infant SD's. Note that for aesthetic reasons they are plotted on the logarithmic scale. The results are striking; across each experimental condition, there is considerable variance heterogeneity at the infant level. Indeed, the proportion of infants that differed from the average within-infant SD (the dotted line) ranged from 0.48 (valid-cue) to 0.63 (baseline). As described in Williams, Martin, and Rast (2019), the average within-infant variance is used to compute traditional reliability indices (Arnon, 2018). Hence, for these data, it is clear that computing reliability based on the average only would mask the large amount of individual differences. And it is not entirely clear, from our perspective, what a single reliability score would mean under the presence of such variation. For example, the maximum–minimum ratio between within-infant variances ranged from 23 (baseline)
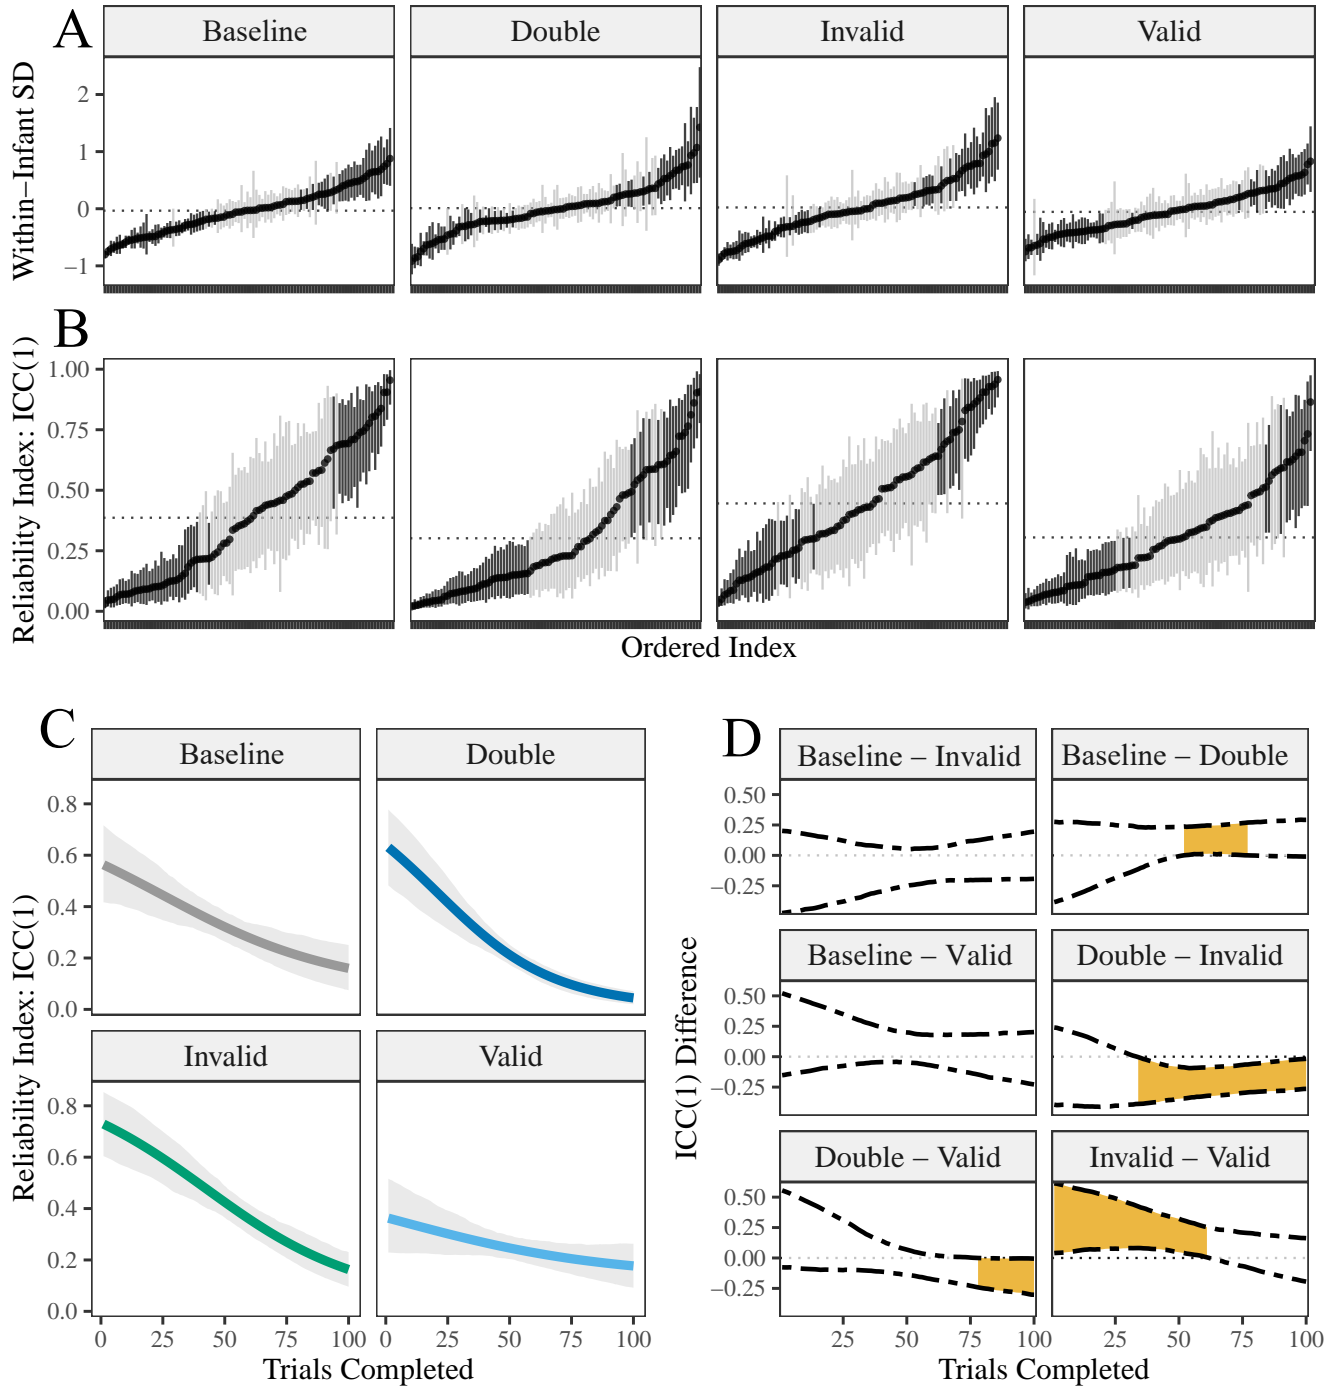
Figure 2. Results for Case 2: Developmental Psychology. A) Within-infant standard deviations (SD) on the log scale. The error bars are 90% credible intervals (CrI). The black error bars excluded the dotted line that denotes the average, or fixed effect, within-infant SD. B) Infant-specific reliability that are separated according to experimental condition. ICC(1) is for single observations. The dotted lines denote the average reliability for each condition. The black error bars excluded the dotted line that denotes the average, or fixed effect, infant-specific reliability. C) Predicted reliability as a function of the number of trials completed. The shaded bands correspond to the posterior SD of the predicted reliability. This is analogous to one standard error. D) Predicted reliability differences between conditions. The dotted line is the "null" value of zero. The shaded regions correspond to areas in which the difference in predicted reliability excluded zero for that value of the covariate (trials completed). The confidence bands correspond to 90% credible intervals.

to 108 (double cue). These ratios were obtained from the shrunken estimates and thus the empirical ratios were much larger. This is further described below (Section *Reliability Indices*).

"True" Score Variance. We also predicted true score variance with the number of trials completed for each infant. These results are included in the Appendix (Table A1). Notably, in all four conditions there was a negative effect, such that, given an *increase* in the number of trials completed, true score variance is expected to *decrease*. In other words, those infants that completed the most trials tended to also be the most homogeneous in the sample. Importantly, the posterior probability of a negative effect exceeded 95% for three conditions. The exception was the invalid cue condition, where the posterior probability was 87%. The double-cue condition in particular had a large effect. For example, an additional 25 trials completed resulted in true score variance reducing from 0.32 to 0.21–a decrease of 34%.

We emphasize that this result does *not* imply, in general, that increasing the number of trials completed reduces true score variance. And we encourage readers to *not* extend this finding beyond what is an explicitly exploratory context. However, this example does make clear that questions related to improving reliability can be investigated at the level of true score variance. This has far reaching implications: true score variance can be a function of covariates, which opens the door for explaining a key aspect of reliability.

Reliability Indices. Figure 2 (panel B) includes the infant-specific estimates of reliability. These were computed with Equation (20), which considers the number of trials completed for each infant and their within-infant variance estimate, respectively. The utility of computing group-specific reliability is clear, as there is substantial variation in reliability at the infant level. Recall this reliability index is ICC(1), which is the expected correlation between any two observations from the same infant. Hence, for many infants, their latencies were highly correlated, whereas, for other infants, their latencies were nearly independent. This insight was made possible by extending CTT to accommodate heterogeneous variance.

It is informative to consider ICC(2), for average scores, which can be computed with Equation (12). Given ICC(1) values of 0.1 and 0.7, and wanting to know the average score reliability of, say, 5 observations, this would translate into ICC(2) values of 0.35 and 0.92. Hence, for some infants, few observations would be needed to reach adequate reliability and other infants would require more measurements.

Panel C includes estimates of reliability that are expressed as a function of the number of trials completed. This index was given in Equation (21). In this example, we predicted ICC(1) with a sequence of trials completed that ranged from 1 to 100. A consistent pattern emerged, in that, for all conditions, reliability is expected to decrease with more trials

completed. This suggests that observations from the same infant approached independence with more trials completed. Said another way, *in these data*, individual differences tend to reduce with completing more trials. This does not imply that gathering many trials is necessarily disadvantageous for detecting individual differences. For example, with few measurements, the individual effects will also have wider credible intervals that effective making it very difficult to differentiate among infants. Hence we encourage researchers to interpret this result with care and consider that true score variance *could* be related to the number of trials completed in this particular task and age group. The presented methodology not only raised this possibility, but it can also be used to look further into this intriguing result.

Panel D also includes predicted reliability as a function of trials completed, but, in this case, it is expressed as a difference between conditions. That is, we predicted reliability given a number of trial completed, as in panel C, and then computed pairwise differences in ICC(1). The shaded regions correspond to values of the covariate (trials completed) that excluded zero. Consider the bottom-right panel, which includes the contrast between the invalid and valid-cue. Here the invalid-cue had notably higher reliability when fewer trials were completed and gradually became similar to the valid-cue condition with more trials completed. This difference was not small, in that, for 25 trials completed, the invalid-cue ICC(1) was 0.60 and the valid-cue ICC(1) was 0.30. On the other hand, while this method does not provide evidence for invariant reliability (i.e., the null hypothesis),[2] it is still informative to note conditions that did not differ from one another. For example, at no value of the covariate (trials completed) did reliability differ between the baseline compared to both the valid and invalid-cue. We emphasize the novelty and utility of this approach. Because we modeled the variance components, this opened the door for investigating reliability with an unprecedented level of detail.

Discussion

In this work, we proposed a novel approach for investigating measurement reliability in hierarchical models. The primary motivation for developing this methodology was that classical test theory appears to be incompatible with commonly observed phenomena in psychological applications. Namely, that homogeneous variance appears to be the exception and not the rule. We noted that homogeneous variance components are assumed to be the case in ANOVA and traditional hierarchical models, and thus also assumed when computing reliability. Our methodology not only relaxes this assumption, but it allows the true score and error variance

---

[2]This could be directly tested, but would require introducing an extension to allow for Bayesian hypothesis testing. This is beyond the scope of the current paper.

to be a function of covariates, including individual variation therein. Hence, the reliability of a questionnaire, measurement device, or experimental task, can be *explained*. This can provide unique insights into reliability that are made possible by merging ideas stemming from the variance modeling literature with classical test theory.

Practical Implications

There are several practical implications of this work. First, in educational research, we demonstrated the utility of modeling the variance structure for assessing standardized test scores. In particular, we demonstrated that both the true score and error variance can be the target of an explanatory model. This can inform study designs, for example sampling strategies, with the goal of utilizing school-level information to obtain a specific level of reliability.

Second, we also described that average score reliability, or ICC(2), is equivalent to the "shrinkage" factor in a one-way random effects model, but it assumes a common variance that can mask important information. Further, the Spearman-Brown prophecy in Equation (12), that predicts reliability given a number of measurements, also assumes a common variance. Our contribution here was to extend both to accommodate heterogeneous variance components. This is especially important because these indices now reflect the substantial variance heterogeneity that is inherent to repeated measurements in psychological science.

Third, in cognitive psychology, there is a debate surrounding the adequacy of commonly used experimental tasks for studying individual differences. The emerging consensus is that for many measures reliability is too low (i.e., "noisy" measures) to adequately study individual variation. However, to our knowledge, this discussion has revolved almost exclusively around the mean structure (i.e., the experimental effect) and avoided the within-person variance structure altogether. This is unfortunate, because reliability has routinely been computed from tasks that have heterogeneous error variance. From our perspective, a satisfactory answer to the question of individual differences in cognitive tasks would require addressing the extreme heterogeneity in within-person variance and thus reliability (e.g., Figure 2, panel A).

Moreover, recall that this work is partially motivated by findings demonstrating that homogeneous within-person variance appears to be the exception and not the norm in psychological applications, including cognitive tasks. This is not readily accounted for by current measurement theories. However, these insights have thus far been restricted to college age students Williams, Martin, and Rast (2019) and older adults Williams, Liu, et al. (2019). In other words, it was not clear whether this general pattern would generalize to within-infant variance. Hence, this work adds to the growing literature on within-person variance heterogeneity that

now spans from infancy to adulthood. And these findings lend further credence to the notion that heterogeneous variance (and thus reliability) is a defining feature of commonly used cognitive tasks. Importantly, we have done more than highlight this issue as it relates to reliability. We have provided an innovate approach for investigating reliability that borrows ideas from the variance modeling literature (Cleveland et al., 2003; Hedeker et al., 2008, 2012; Lindley, 1969).

Methodological Implications

The primary aim of this work was to provide the necessary ingredients to extend classical test theory. Of course, various alternatives have been proposed, including generalizability theory (Brennan, 1992) and item response theory (DeMars, 2018). Our aim, however, was to explicitly stay within the CTT framework, but with the addition of modeling true score and error variance. Accordingly, the notion of how observed scores come about remain in tact, but with the possibility of predicting the variance components. Hence, rather than viewing reliability as a stable property of, for example, an experimental task, we can fit explanatory models with the goal of finding variables or situations that influence (positively) reliability. Essentially, by identifying predictors of true score or error variance, we can refine measurement. Such possibilities were demonstrated in the case studies, although these examples just scratched the surface of possibilities. For example, we focused on predicting true score variance, whereas identifying sources of within-group variance is also important to consider (i.e., error variance, Karch et al., 2019)

Historical context should be taken into account when considering to expand a theory. CTT (and *G*-theory) is inextricably linked to ANOVA (Fisher, 1925), that is, "Often, CTT and analysis of variance (ANOVA) are viewed as the parents of *G*-theory" (p. 7, Brennan, 2011), with its origins going back to a time when the primary mode of transportation was horse and buggy. Just as the horse can limit possibilities for travel, so too can ANOVA (and traditional mixed models) for exploring measurement reliability. In this respect, our work can be seen as an extension to CTT; that is, a natural progression from classical to modern methodology. This brings full circle the arguments of Lindley (1969) and Novick et al. (1971), where it was originally noted that Bayesian methods offer flexibility not possible with ANOVA.[3] The user-friendly R package **ICCier** can facilitate this transition towards richer models for investigating measurement reliability, all the while staying within the CTT framework.

---

[3]The R package `nlme` can be used for simple heterogeneous variance models. There are key disadvantages, however, including that the variances are not hierarchically modeled and that a measure of uncertainty is not readily available. Further, the models fitted in this paper failed to converge in `nlme`.

An Alternative Perspective

It would be remiss of us to not offer an alternative perspective. It is customary to view the residuals as mere "noise" and perhaps measurement "error." For example, trial to trial fluctuations are a nuisance to understanding the latent process. On the other hand, there is a large body of literature that views these same fluctuations as a key aspect of the construct. A good example is personality traits, that were typically considered fixed, but now an active area of research revolves around within-person variability of these traits (i.e., the fluctuations; Fleeson, 2001; Hutteman, Back, Geukes, Küfner, & Nestler, 2016; Williams, Liu, et al., 2019). So rather than there being individual differences in reliability, the alternative perspective is to view these as individual differences in stability. That is, individuals with larger residual variance are relatively more volatile or inconsistent. Further, instead of using an adjective to label between-group variance "true," this perspective would view homogeneity of means as an important aspect of the construct. Consider age groups. An interesting research question is whether older adults are more or less homogeneous in, for example, cognitive abilities. Such a conclusion is diametrically opposed to CTT, and the reliability literature in general, where measurements are construed as a true score plus error. However, in our opinion, this alternative perspective is worth considering.

Future Directions

The proposed methodology provides a foundation for further quantitative advances. First, although we stayed within the CTT framework, the basic ideas readily apply to *G*-theory. In this case, explanatory models could be fitted to several variance components. The model described in the section Model Formulation can seamlessly accommodate this extension. Second, this methodology could be extended to consider test-retest reliability. This is especially important, because, to our knowledge, heterogeneous within-group variance has not been considered in test-retest situations.

Conclusion

Measurement reliability has traditionally been considered a stable property of a measurement device or experimental task. This framework does not allow for *explaining* reliability, because it assumes the true score and error variance are fixed and non-varying. However, heterogeneous variance modeling allows for investigating reliability with a fine-tooth comb. The illustrative examples highlighted such possibilities. There can be substantial variability in error variance, which necessarily implies the same for reliability, and that true score variance can be a function of covariates. The methodology that made these insights possible is implemented in the R package **ICCier**.

References

Aitkin, M. (1987, 6). Modelling Variance Heterogeneity in Normal Regression Using GLIM. *Applied Statistics*, *36*(3), 332. doi: 10.2307/2347792

Arnon, I. (2018). Do current statistical learning capture stable individual differences in children? An investigation of task reliability across modalities.

Bartko, J. J. (1966, 8). The Intraclass Correlation Coefficient as a Measure of Reliability. *Psychological Reports*, *19*(1), 3–11. doi: 10.2466/pr0.1966.19.1.3

Betancourt, M. (2017, 1). A Conceptual Introduction to Hamiltonian Monte Carlo. *arXiv*.

Bliese, P. D. (2000). *Within-group agreement, non-independence, and reliability: Implications for data aggregation and analysis* (1st ed.; K. J. Klein & S. W. Kozlowski, Eds.). San Francisco: Jossey-Bass.

Brennan, R. L. (1992). Generalizability theory. *Educational Measurement: Issues and Practice*, *11*(4), 27–34.

Brennan, R. L. (2011). Generalizability theory and classical test theory. *Applied Measurement in Education*, *24*(1), 1–21. doi: 10.1080/08957347.2011.532417

Brown, W. (1910). SOME EXPERIMENTAL RESULTS IN THE CORRELATION OF MENTAL ABILITIES. *British Journal of Psychology, 1904 [U+2010] 1920*, *3*(3), 296–322. doi: 10.1111/j.2044-8295.1910.tb00207.x

Brunton-Smith, I., Sturgis, P., & Leckie, G. (2017). Detecting and understanding interviewer effects on survey data by using a cross-classified mixed effects location–scale model. *Journal of the Royal Statistical Society. Series A: Statistics in Society*, *180*(2), 551–568. doi: 10.1111/rssa.12205

Cleveland, W. S., Denby, L., & Liu, C. (2003). Random scale effects. (2), 33. Retrieved from `stat.bell-labs.com`

Colombo, J., & Fagen, J. (2014). *Individual differences in infancy: Reliability, stability, and prediction*. Psychology Press.

Cronbach, L. J. (1972). The dependability of behavioral measurements. *Theory of generalizability for scores and profiles*, 1–33.

de Boer, H., Timmermans, A. C., & van der Werf, M. P. C. (2018, 4). The effects of teacher expectation interventions on teachers' expectations and student achievement: narrative review and meta-analysis. *Educational Research and Evaluation*, *24*(3-5), 180–200. doi: 10.1080/13803611.2018.1550834

DeMars, C. E. (2018, 2). Classical Test Theory and Item Response Theory. In *The wiley handbook of psychometric testing* (pp. 49–73). Chichester, UK: John Wiley & Sons, Ltd. doi: 10.1002/9781118489772.ch2

Efron, B., & Morris, C. (1977). Stein's paradox in statistics. *Scientific American*, *236*(5), 119–127.

Fisher, R. A. (1925). *Statistical methods for research workers*. London: Oliver & Bond.

Fleeson, W. (2001). Toward a structure-and process-integrated view of personality: Traits as density distributions of states. *Journal of personality and social psychology*, *80*(6), 1011–27.

Gelfand, A., & Smith, A. F. M. (1990). Sampling-Based Approaches to Calculating Marginal Densities. *Journal Of The American Statistical Association*, *85*(410), 398–409. doi: 10.2307/2289776

Gelman, A., Hill, J., & Yajima, M. (2012, 4). Why We (Usually)

Don't Have to Worry About Multiple Comparisons. *Journal of Research on Educational Effectiveness*, *5*(2), 189–211. doi: 10.1080/19345747.2011.618213

Goldstein, H., Rasbash, J., Yang, M., Woodhouse, G., Pan, H., Nuttall, D., & Thomas, S. (1993). A Multilevel Analysis of School Examination Results. *Oxford Review of Education*, *19*(4), 425–433. doi: 10.1080/0305498930190401

Gu, Z., Emons, W. H. M., & Sijtsma, K. (2018). Review of issues about classical change scores: A multilevel modeling perspective on some enduring beliefs. *psychometrika*, *83*(3), 674–695.

Hamaker, E. L., Asparouhov, T., Brose, A., Schmiedek, F., & Muthén, B. (2018). At the Frontiers of Modeling Intensive Longitudinal Data: Dynamic Structural Equation Models for the Affective Measurements from the COGITO Study. *Multivariate Behavioral Research*, *53*(6), 820–841. doi: 10.1080/00273171.2018.1446819

Haveman, R., & Smeeding, T. (2006). The role of higher education in social mobility. *The Future of children*, *16*(2), 125–50.

Hedeker, D., Mermelstein, R. J., & Demirtas, H. (2008). An application of a mixed-effects location scale model for analysis of ecological momentary assessment (EMA) data. *Biometrics*, *64*(2), 627–634. doi: 10.1111/j.1541-0420.2007.00924.x

Hedeker, D., Mermelstein, R. J., & Demirtas, H. (2012). Modeling between-subject and within-subject variances in ecological momentary assessment data using mixed-effects location scale models. *Statistics in Medicine*, *31*(27), 3328–3336. doi: 10.1002/sim.5338

Hedge, C., Powell, G., & Sumner, P. (2018). The reliability paradox: Why robust cognitive tasks do not produce reliable individual differences. *Behavior Research Methods*, *50*(3), 1166–1186. doi: 10.3758/s13428-017-0935-1

Hoffman, L. (2007). Multilevel Models for Examining Individual Differences in Within-Person Variation and Covariation Over Time. *Multivariate Behavioral Research*, *42*(4), 609–629. doi: 10.1080/00273170701710072

Hutteman, R., Back, M. D., Geukes, K., Küfner, A. C., & Nestler, S. (2016). Trait personality and state variability: Predicting individual differences in within- and cross-context fluctuations in affect, self-evaluations, and behavior in everyday life. *Journal of Research in Personality*, *69*, 124–138. doi: 10.1016/j.jrp.2016.06.003

Karch, J. D., Filevich, E., Wenger, E., Lisofsky, N., Becker, M., Butler, O., ... Kühn, S. (2019, 10). Identifying predictors of within-person variance in MRI-based brain volume estimates. *NeuroImage*, *200*, 575–589. doi: 10.1016/J.NEUROIMAGE.2019.05.030

Kelley, T. L. (1927). *The Interpretation of Educational Measurements*. New York: World Book.

Kline, T. (2005). *Psychological testing: A practical approach to design and evaluation*. Thousand Oaks: Sage.

Leckie, G., French, R., Charlton, C., & Browne, W. (2014). Modeling heterogeneous variance–covariance components in two-level models. *Journal of Educational and Behavioral Statistics*, *39*(5), 307–332.

Lindley, D. (1969, 12). A BAYESIAN ESTIMATE OF TRUE SCORE THAT INCORPORATES PRIOR INFORMATION. *ETS Research Bulletin Series*, *1969*(2), i-8. doi: 10.1002/j.2333-8504.1969.tb00754.x

Logan, J. R., & Burdick-Will, J. (2017, 11). School Segregation and Disparities in Urban, Suburban, and Rural Areas. *The Annals of the American Academy of Political and Social Science*, *674*(1), 199. doi: 10.1177/0002716217733936

Lord, F. M., & Novick, M. R. (1968a). *Statistical theories of mental test scores*. Oxford: IAP.

Lord, F. M., & Novick, M. R. (1968b). *Statistical theories of mental test scores*. Reading: Adison-Wesley.

McCulloch, C. E. (2003). Chapter 2: Linear mixed models (LMMs). In *Generalized linear mixed models* (Vol. Volume 7, pp. 9–20). Beechwood OH and Alexandria VA: Institute of Mathematical Statistics and American Statistical Association.

McGraw, K. O., & Wong, S. P. (1996). Forming inferences about some intraclass correlation coefficients. *Psychological Methods*, *1*(1), 30–46. doi: 10.1037/1082-989X.1.1.30

Morris, C. N., & Lysy, M. (2012). Shrinkage Estimation in Multilevel Normal Models. *Statistical Science*, *27*(1), 115–134. doi: 10.1214/11-sts363

Myles, J., Price, G., Hunter, N., Day, M., & Duffy, S. (2003). A potentially useful distribution model for dietary intake data. *Public Health Nutrition*, *6*(5), 513–519. doi: 10.1079/phn2003459

Novick, M. R. (1965). The axioms and principal results of classical test theory. *ETS Research Bulletin Series*, *1965*(1), i–31.

Novick, M. R., Jackson, P. H., & Thayer, D. T. (1971, 9). Bayesian inference and the classical test theory model: Reliability and true scores. *Psychometrika*, *36*(3), 261–288. doi: 10.1007/BF02297848

Rabe-Hesketh, S., & Skrondal, A. (2008). *Multilevel and longitudinal modeling using Stata*. STATA press.

Rast, P., & Ferrer, E. (2018). A Mixed-Effects Location Scale Model for Dyadic Interactions. , 1–63. doi: 10.1080/00273171.2018.1477577

Rast, P., Hofer, S. M., & Sparks, C. (2012). Modeling Individual Differences in Within-Person Variation of Negative and Positive Affect in a Mixed Effects Location Scale Model Using BUGS/JAGS. *Multivariate Behavioral Research*, *47*(2), 177–200. doi: 10.1080/00273171.2012.658328

Raudenbush, S. W., & Bryk, A. S. (2002). *Hierarchical Linear Models: Applications and Data Analysis Methods* (2nd ed.). Thousand Oaks: Sage Publications.

Robert, C., & Casella, G. (2011, 8). A Short History of Markov Chain Monte Carlo: Subjective Recollections from Incomplete Data. *Statistical Science*, *26*(1), 102–115. doi: 10.1214/10-STS351

Rose, S. A., Feldman, J. F., & Wallace, I. F. (1988). Individual differences in infants' information processing: Reliability, stability, and prediction. *Child Development*, 1177–1197.

Ross-Sheehy, S., Schneegans, S., & Spencer, J. P. (2015). The infant orienting with attention task: Assessing the neural basis of spatial attention in infancy. *Infancy*, *20*(5), 467–506.

Rouder, J. N., Kumar, A., & Haaf, J. M. (2019). Why most studies of individual differences with inhibition tasks are bound to fail. *PsyArXiv*, 1–37.

Sandy, J., & Duncan, K. (2010, 9). Examining the achievement test score gap between urban and suburban students.

Education Economics, *18*(3), 297–315. doi: 10.1080/09645290903465713

Shavelson, R. J., & Webb, N. M. (2010). Generalizability Theory. In J. L. Green, G. Camilli, & P. B. Elmore (Eds.), *Handbook of complementary methods in education research* (Vol. 24, pp. 27–34). Thousand Oaks: Wiley Online Library.

Shavelson, R. J., & Webb, N. M. (2012). Generalizability Theory. In J. Green, G. Camilli, & P. Elmore (Eds.), *Handbook of complementary methods in education research* (3rd ed.). Routledge.

Shieh, G. (2016). Choosing the best index for the average score intraclass correlation coefficient. *Behavior Research Methods*, *48*(3), 994–1003. doi: 10.3758/s13428-015-0623-y

Sirin, S. R. (2005, 9). Socioeconomic Status and Academic Achievement: A Meta-Analytic Review of Research. *Review of Educational Research*, *75*(3), 417–453. doi: 10.3102/00346543075003417

Spearman, C. (1910, 10). CORRELATION CALCULATED FROM FAULTY DATA. *British Journal of Psychology, 1904-1920*, *3*(3), 271–295. doi: 10.1111/j.2044-8295.1910.tb00206.x

Stan Development Team. (2016). *Rstan: the R interface to Stan.* Retrieved from http://mc-stan.org/

Stein, C. (1956). Inadmissibility of the usual estimator for the mean of a multivariate normal distribution. In *Proceedings of the {t}hird {b}erkeley {s}ymposium on {m}athematical {s}tatistics and {p}robability, 1954–1955, vol. {i}* (pp. 197–206). University of California Press, Berkeley and Los Angeles.

Thorndike, E. L. (1904). *An introduction to the theory of mental and social measurements*. The Science Press.

Thouless, R. H. (1936). Test Unrealiability and Function Fluctuation. *British Journal of Psychology*, *26*(4), 325.

Traub, R. E. (1997). Classical test theory in historical perspective. *Educational Measurement*, *16*, 8–13.

Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, *27*(5), 1413–1432. doi: 10.1007/s11222-016-9696-4

Watts, A., Walters, R. W., Hoffman, L., & Templin, J. (2016). Intra-individual variability of physical activity in older adults with and without mild Alzheimer's disease. *PLoS ONE*, *11*(4). doi: 10.1371/journal.pone.0153898

Williams, D. R., Liu, S., Martin, S. R., & Rast, P. (2019). Bayesian Multivariate Mixed-Effects Location Scale Modeling of Longitudinal Relations among Affective Traits, States, and Physical Activity. doi: 10.31234/OSF.IO/4KFJP

Williams, D. R., Martin, S. R., & Rast, P. (2019). Putting the Individual into Reliability: Bayesian Testing of Homogeneous Within-Person Variance in Hierarchical Models. *PsyArXiv*. doi: 10.31234/OSF.IO/HPQ7W

Williams, D. R., Zimprich, D. R., & Rast, P. (2019, 5). A Bayesian nonlinear mixed-effects location scale model for learning. *Behavior Research Methods*, 1–19. doi: 10.3758/s13428-019-01255-9

Appendix A
Supplementary Results

Table A1

Sub-model estimates for predicting "true" score variance

| | Estimates | | | |
Parameter	M	SD	90% CrI	$p(\kappa_1 < 0|\mathbf{Y})$
$\exp[\kappa_0^{(control)}]$	0.37	0.05	[0.31, 0.45]	–
$\exp[\kappa_0 + (\kappa_1 \cdot 2.5)]$	0.30	0.06	[0.21, 0.41]	0.97
$\exp[\kappa_0^{(double)}]$	0.32	0.05	[0.26, 0.39]	–
$\exp[\kappa_0 + (\kappa_1 \cdot 2.5)]$	0.21	0.06	[0.14, 0.29]	0.99
$\exp[\kappa_0^{(invalid)}]$	0.45	0.05	[0.38, 0.53]	–
$\exp[\kappa_0 + (\kappa_1 \cdot 2.5)]$	0.35	0.07	[0.27, 0.46]	0.99
$\kappa_0^{(valid)}$	0.38	0.05	[0.31, 0.46]	–
$\exp[\kappa_0 + (\kappa_1 \cdot 2.5)]$	0.34	0.07	[0.25, 0.45]	0.87

Note. Posterior mean (M) and standard deviation (SD). κ_0 is the intercept and the true score variance for the average number of completed trials. κ_1 is the regression weight. $p(\kappa_1 < 0|\mathbf{Y})$ denotes the posterior probability of a negative effect. To ease interpretation, we computed the true score variance for an increase of 25 trials completed. This corresponds to $\exp[\kappa_0 + (\kappa_1 \cdot 2.5)]$.

Appendix B
Example Code

All analyses were obtained from the **ICCier** package (development version: https://github.com/stephensrmmartin/ICCier). This version can be installed into R via devtools::install_github('stephensrmmartin/ICCier'). Once installed, we can load the ICCier package and the school-data from the mlmRev package:

```
library(ICCier)
library(mlmRev)
school_dat = mlmRev::Exam

iccier = ICCier(x = normexam,
                between ~ vr,
                group = school,
                data = school_dat)
```

The ICCier function takes the independent variable in x (here, normexam). The between argument defines the submodel for the between-group variance. Here, vr is defines the schools intake level quartiles. The grouping variable is defined in group.

A summary of the fitted object can be obtained via

```
> summary(iccier)
Formula: normexam | school ~ 1 | 1 | vr
Type:  Unconditional
Number of observations: 4059
Number of groups: 65

--------------------
```

```
Diagnostics:
        Rhats: Passed
        Divergent transitions: Passed

--------------------
ICC Summary:
       Estimate  2.5% 97.5%
Mean     0.171 0.122 0.230
SD       0.105 0.054 0.172


--------------------
Coefficients:

Mean Model:
            Estimate   2.5% 97.5%
(Intercept)   -0.058 -0.161 0.045


Within-group (log) SD:
            Estimate   2.5%  97.5%

(Intercept)   -0.108 -0.153 -0.061


Between-group (log) SD:
Estimate   2.5%  97.5%
Mean_(Intercept)                -0.785 -1.210 -0.324
Mean_(Intercept):vrmid 50%   -0.494 -1.027  0.042
Mean_(Intercept):vrtop 25%    0.237 -0.415  0.867

Random Effect Correlations:
Mean_(Intercept) (Intercept)
Mean_(Intercept)              1.000       0.315
(Intercept)                  0.315       1.000
```