

UC Davis

UC Davis Previously Published Works

Title

Beneath the surface: Unearthing within-person variability and mean relations with Bayesian mixed models.

Permalink

<https://escholarship.org/uc/item/8j6856rp>

Journal

Psychological Methods, 26(1)

ISSN

1082-989X

Authors

Williams, Donald R
Mulder, Joris
Rouder, Jeffrey N
[et al.](#)

Publication Date

2021-02-01

DOI

10.1037/met0000270

Peer reviewed



HHS Public Access

Author manuscript

Psychol Methods. Author manuscript; available in PMC 2021 November 07.

Published in final edited form as:

Psychol Methods. 2021 February ; 26(1): 74–89. doi:10.1037/met0000270.

Beneath the Surface: Unearthing Within-Person Variability and Mean Relations with Bayesian Mixed Models

Donald R. Williams,
University of California, Davis

Joris Mulder,
Tilburg University

Jeffrey N. Rouder,
University of California, Irvine

Philippe Rast
University of California, Davis

Abstract

Mixed-effects models are becoming common in psychological science. Although they have many desirable features, there is still untapped potential. It is customary to view homogeneous variance as an assumption to satisfy. We argue to move beyond that perspective, and to view modeling within-person variance as an opportunity to gain a richer understanding of psychological processes. The technique to do so is based on the mixed-effects location scale model that can simultaneously estimate mixed-effects sub-models to both the mean (location) and within-person variance (scale). We develop a framework that goes beyond assessing the sub-models in isolation of one another and introduce a novel Bayesian hypothesis test for mean-variance correlations in the distribution of random effects. We first present a motivating example, which makes clear how the model can characterize mean–variance relations. We then apply the method to reaction times gathered from two cognitive inhibition tasks. We find there are more individual differences in the within-person variance than the mean structure, as well as a complex web of structural mean–variance relations. This stands in contrast to the dominant view of within-person variance (i.e., “noise”). The results also point towards paradoxical within-person, as opposed to between-person, effects: several people had *slower* and *less* variable incongruent responses. This contradicts the typical pattern, wherein *larger* means tend to be associated with *more* variability. We conclude with future directions, spanning from methodological to theoretical inquiries, that can be answered with the presented methodology.

Keywords

Bayesian; mean–variance relations; mixed-effects location scale model; Bayes factor; within-person variance; behavioral (in)consistency

Introduction

Repeated measurement designs are common to the social-behavioral sciences. Their use spans from observational inquiries that track individuals over an extended period of time, to controlled settings that can include hundreds of experimental trials for each person. Modeling these kinds of data requires techniques that are able to partition and account for different sources of variation, for example, in the experimental effect (Aarts, Dolan, Verhage, & van der Sluis, 2015) or stimulus type (Wolsiefer, Westfall, & Judd, 2017). Adequately accounting for these sources of variability leads to the desired (frequentist) inference by ensuring that nominal error rates are maintained (Aarts, Verhage, Veenvliet, Dolan, & van der Sluis, 2014; Barr, Levy, Scheepers, & Tily, 2013; Judd, Westfall, & Kenny, 2012; Williams, Carlsson, & Bürkner, 2017). The idea here is that researchers need to control or correct for variability. On the other hand, modeling these same sources of variation can provide valuable insight into psychological processes that drive the variation itself. A prominent example is the study of individual differences in, say, temporal changes (Liu, Rovine, & Molenaar, 2012), inhibition (Haaf & Rouder, 2017), or learning trajectories (Williams & Rast, 2018).

Although there are many approaches to modeling repeated measurements data, two general trends have emerged in psychology. First, some researchers seek to construct elaborate, nonlinear models, that are meant to mimic, as closely as possible, the latent data-generating mechanisms (Mazur, 2006). The resulting models are termed *process models*, and their study has been an integral part of the field of mathematical psychology (see references provided in: Townsend, 2008). The second trend is to focus exclusively on the mean structure—that is, how the mean of the outcome (e.g., response times) varies with, perhaps, an experimental manipulation. In these situations, it is common to aggregate repeated measures data at the individual level (Davidson, Zacks, & Williams, 2003; Wright, 2017), such that each person contributes only their respective mean score. This allows for using a relatively simple statistical method such as the dependent samples *t*-test. Despite their popularity, these approaches are not without limitations (Leppink, 2019; Leppink & Merriënboer, 2015), including removal of the within-person variance and they cannot provide information about individual variability (Bauer, 2011). More recent approaches employ mixed-effect models that can simultaneously estimate the average effect across individuals and the person-specific estimates (i.e., random effects). Their inherent flexibility allows for investigating research questions beyond the average, for example, whether all people had an effect in the predicted direction (Haaf & Rouder, 2017). These developments have been integral in the study of individual differences.

In this work, we argue there is a fertile middle ground between positing full-blown process models and studying mean structures. For the study of individual differences in particular, there is untapped potential in mixed-effects models. The focus has mostly remained on the mean structure, that is the dependent variable of interest, whereas within-person variability is often relegated to “error” and considered a nuisance. In other words, the unexplained variance, or the within-person variance, goes into the residual component where it is typically considered a fixed, non-varying constant, in mixed-effects models. Herein lies the untapped potential: rather than viewing homogeneous variance as an assumption to satisfy,

or within-person variance as “noise,” we can seek to understand it just like the dependent variable.

Despite that in psychology the primary focus has been on the mean structure, there is an interesting and storied literature on modeling within-person variance (i.e, the residual variance, Cleveland, Denby, & Liu, 2003; Hedeker, Mermelstein, & Demirtas, 2008, 2012; Leckie, French, Charlton, & Browne, 2014; Rast & Ferrer, 2018; Rast, Hofer, & Sparks, 2012a). In applied settings, residual variability is termed *intraindividual* variability (IIV; Christ, Combrinck, & Thomas, 2018; Fagot et al., 2018; Röcke & Brose, 2013), and thought to reflect behavioral consistency (Rast, Hofer, & Sparks, 2012b) or predictability (D. J. Mitchell, Fanson, Beckmann, & Biro, 2016). This conceptualization builds upon a central idea that within-person variance, or IIV, is not regarded as reflecting mere measurement error but conveys systematic information (Cattell, Cattell, & Rhymer, 1947; Fiske & Rice, 1955; Horn, 1972; Ram & Gerstorff, 2009; Woodrow, 1932). As an example, consistency in cognitive abilities has been proposed as an indicator of Alzheimer’s disease in aging populations (Kalin et al., 2014). And inconsistency was even suggested to predict death (MacDonald, Hulstsch, & Dixon, 2008). That is, those that were more variable tended to die before those that were relatively stable.

Furthermore, it has been noted that behavioral measures may have important signatures in both the mean and variance structures. For example, Luce (1986) and Wagenmakers and Brown (2007) proposed a lawful relationship in response time: the standard deviation of response time tends to increase linearly with the mean. Rouder and colleagues build in this lawfulness explicitly by using shift-scale-shape models (e.g. Rouder, Tuerlinckx, Speckman, Lu, & Gomez, 2008; Rouder, Yue, Speckman, Pratte, & Province, 2010), such as a three-parameter lognormal with effects in scale. Moreover, common graphical approaches, such as QQ plots and the delta plot (De Jong, Liang, & Lauber, 1994), are commonly employed techniques for visualizing the mean-variance relationship (Ridderinkhof, Scheres, Oosterlaan, & Sergeant, 2005; Schwarz & Miller, 2012). These inquiries into the mean-variance relation share a common thread, that is, they are typically descriptive in nature. And the modeling based approaches, as in Rouder et al. (2010), do not directly target the within-person variance.

The aim of this work is to present a novel, formal Bayesian modeling framework, that allows for testing and visualizing what has been previously stipulated about the mean-variance relation. But, importantly, we focus on *within-person* variance or IIV in hierarchical models. This places our work within the tradition of individual difference research, in that, along with the means, person-specific consistency is characterized. This necessarily leads to rich inferences relating to mean-variance lawfulness, because the correlations between individual difference parameters in the random effects distribution can be tested. Hence, with an uncanny level of detail, researchers can begin to test lawful relations between, say, an experimental effect on reaction time and reaction time consistency.

This work is organized as follows. In the first section we introduce a relatively simple model, where it is made clear how our approach can capture and test mean-variance relations. Additionally, we introduce our hypothesis testing strategy for the correlations

that capture the mean–variance relations in the distribution of random effects. We then apply the methodology to two well-known cognitive inhibition tasks. Here, general advantages of our approach are highlighted, as compared to a traditional mixed-effects model, and the hypothesis testing strategy is employed. We conclude by summarizing our major contribution, limitations of the mean-variance modeling in general, and specific future directions for psychological applications.

Motivating Example

The foundation for our methodology merges two disparate lines of research. On the one hand, the heterogeneous variance modeling literature that attempts to explain within-person variance (Cleveland et al., 2003; J. Foulley & Quaas, 1995; J. L. Foulley, San Cristobal, Gianola, & Im, 1992). Central to this endeavour is the recently proposed mixed-effects location scale model (MELSM, pronounced mel-zem; Hedeker et al., 2008), described below, which allows the within-person variance to be a function of its own mixed-effects model (e.g., Watts, Walters, Hoffman, & Templin, 2016; Williams & Rast, 2018). On the other hand, for testing mean–variance relations, we draw upon the Bayesian literature that employs mixture prior distributions for variable selection—that is, spike and slab methodology (see references in O’Hara & Sillanpää, 2009). As shown in this motivating example, this opens the door for answering novel research questions about the interplay between the mean structure and within-person variance in psychology.

Random Intercept Model

We use data from an inhibition task that investigated the so-called “Stroop Effect.” These data were first reported in von Bastian, Souza, and Gade (2016). They consist of 121 participants, each of which completed approximately 90 trials in total. About half of the trials were in the congruent condition, wherein the number of characters matched the displayed numbers (e.g., 22). The remaining trials were in the incongruent condition (e.g., 222). Further details are provided below (Section Illustrative Examples). The outcome is reaction time (on the seconds scale) for correctly identifying the number of characters.

Mean Structure.—For the i th person and j th trial, the mean structure is defined as

$$y_{ij} = \beta_0 + u_{0i} + \epsilon_{ij}, \quad (1)$$

where β_0 is the fixed effect, or the population-averaged mean, and u_{0i} is the individual deviation, or random effect, for subject i . The random effects are assumed to follow a normal distribution with mean zero. Hence, for person i , their respective response time mean is $\beta_0 + u_{0i}$. Note this is not equivalent to estimating the empirical means, because the hierarchical structure, described below (Equation 3), provides shrinkage that improves accuracy by smoothing the random effects towards β_0 (Efron & Morris, 1977; James & Stein, 1961).

Variance Structure.—We next account for the “errors”. While they are typically assumed to be normally distributed, with a constant variance, this is not the case for the MELSM, that is,

$$\begin{aligned} \epsilon_{ij} &\sim \mathcal{N}(0, \sigma_{\epsilon_{ij}}^2) \text{ with} \\ \sigma_{\epsilon_{ij}}^2 &= \exp(\eta_0 + u_{1i}). \end{aligned} \quad (2)$$

The subscripts denote the residual for the i th person and j th trial. Further, the error variance $\sigma_{\epsilon_{ij}}^2$ is now allowed to vary across i individuals given a log-linear model (Hedeker et al., 2008). These parameters are analogous to those in Equation (1), in that η_0 is the average of the within-person variances on the logarithmic scale. The random effects u_{1i} are the individual deviations from η_0 and they are assumed to be normally distributed with mean zero (Equation 3). This allows each person to have their own “error” variance. For person i , $\eta_0 + u_{1i}$ quantifies the variability of their respective response time distribution. We assume diffuse prior distributions for the fixed effects, $\beta_0, \eta_0 \sim \mathcal{N}(0, 5)$, which is justifiable because they are not tested in this example.

Random Effects Covariance Matrix.—This work focuses on the relations between the means and within-person variances. As such, we assume that the individual effects from both sub-models are sampled from a common multivariate normal distribution—that is,

$$\begin{bmatrix} u_{0i} \\ u_{1i} \end{bmatrix} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}). \quad (3)$$

with mean zero and $\mathbf{\Sigma}$ is the covariance matrix that includes the random effects variances and covariances. A variety of priors have been proposed for $\mathbf{\Sigma}$ (see references in Alvarez, Niemi, & Simpson, 2014). Historically, the inverse-Wishart distribution has been a popular choice because it is the conjugate prior for $\mathbf{\Sigma}$ (Gutiérrez-Peña et al., 1997). However, it has been criticized for being overly restrictive (e.g., a common parameter governs *all* elements; Hsu, Sinay, & Hsu, 2012; Leonard & Hsu, 1992). Further, recall that our aim is *covariance* selection in the distribution of random effects (Equation 3). Hence, the random effects correlations, which are scale free, are a natural target for Bayesian variable selection. To this end, we use the separation strategy to decompose $\mathbf{\Sigma}$ (see Equation 1 in Barnard, McCulloch, & Meng, 2000). This can be written as

$$\mathbf{\Sigma} = \text{diag}(\boldsymbol{\tau}) \mathbf{\Omega} \text{diag}(\boldsymbol{\tau}), \quad (4)$$

where $\boldsymbol{\tau}$ is a 2×1 vector, $\tau_i = 0, 1$, that contains the random effects standard deviations $S D$, $\text{diag}(\boldsymbol{\tau})$ is a diagonal matrix with the diagonal elements $\boldsymbol{\tau}$, and $\mathbf{\Omega}$ is a 2×2 correlation matrix. Note that τ_i refers to element i in $\boldsymbol{\tau}$. With this decomposition, we can specify independent priors for each element of $\boldsymbol{\tau}$ and $\mathbf{\Omega}$. For the random effects $S D$ s, we again assume diffuse priors,

$$\tau_i \sim \text{Student-t}^+(\mu = 0, \sigma = 1, \nu = 10), \quad i = 0, 1, \quad (5)$$

where each is assigned a half Student-t distribution. This family of priors was proposed in Gelman (2006) and then Huang and Wand (2013) extended the idea to multivariate settings. Note that we are not testing for the presence of variability, for example $\tau_1 > 0$, but instead

the mean–variance relation captured in $\mathbf{\Omega}$. We return to the important topic of testing, say, the fixed effects and random effects SD s, in the discussion.

Mixture Prior Distribution.: We now describe the prior for the correlations in $\mathbf{\Omega}$ (Equation 4). We employ the spike and slab approach for variable selection. In this approach, model comparison is typically formulated as a *two* component mixture: 1) a “spike” that is either narrowly concentrated around zero (George & McCulloch, 1993; George & McCulloch, 1997) or a point mass at zero (Kuo & Mallick, 1998; T. J. Mitchell & Beauchamp, 1988) and 2) a diffuse “slab” component surrounding zero. A central aspect of this approaches is the addition of an indicator variable (Kuo & Mallick, 1998), which in essence allows for switching between the mixture components (i.e., transdimensional MCMC, Heck, Overstall, Gronau, & Wagenmakers, 2018). The proportion of MCMC samples spent in each component can then be used to approximate the respective posterior model probabilities or the marginal Bayes factors.

Because we target the random effects distribution, this requires an innovative approach for *covariance* selection. We build upon approaches described in Wang (2015) and Frühwirth-Schnatter and Tüchler (2008), each of which employed a mixture of continuous distributions. The former was in the context of graphical models, whereas the latter was also for random effects covariance matrices. They used a binary indicator, whereas our innovation is to allow for several competing hypotheses. This is accomplished by introducing a three component mixture prior for the random effects correlations.

Recall that the random effects correlations are the off-diagonal elements of $\mathbf{\Omega}$, that is,

$$\mathbf{\Omega} = \begin{bmatrix} 1 & \rho_{01} \\ \rho_{01} & 1 \end{bmatrix}. \quad (6)$$

We employ the categorical distribution, which generalizes the Bernoulli distribution, to the case of several categories

$$k_{01} \sim \text{Cat}(\boldsymbol{\pi}), k \in \{1, 2, 3\}. \quad (7)$$

Here $\boldsymbol{\pi}$ is a 1×3 vector of prior probabilities for each category, such that $\sum_{k=1}^K \pi_k = 1$, and k_{01} can be understood as an indicator variable. We assume equal prior probabilities $1/K$. It follows that the prior for ρ_{01} , that captures the relations between the means and within-person variances, is then a mixture of three distributions. In our formulation, this mixture is defined as

$$\rho_{01} = F^{-1}(z_{01}) \quad (8)$$

$$z_{01} \sim \begin{cases} \mathcal{N}(0, s_\rho \cdot c_\rho^{-1}) & \text{if } k = 1 \text{ } c_\rho \gg 1 \\ \mathcal{N}^+(0, s_\rho) & \text{if } k = 2 \\ \mathcal{N}^-(0, s_\rho) & \text{if } k = 3, \end{cases}$$

where \mathcal{N}^+ is a half-normal distribution restricted to positive values, \mathcal{N}^- is a half-normal distribution restricted to negative values. Consequently, s_ρ is the scale of a half-normal ($k = 2, 3$) or a normal distribution ($k = 1$), respectively. After taking the inverse of the Fisher z transformation, $F^{-1}(z_{01})$, this results in the prior for ρ_{01} (Equation 6). Assuming a prior on z was described in Daniels and Kass (section 2.3.2, 1999), which was motivated by approaches for covariance matrix estimation (pp. 5 - 6 in Lin & Perlman, 1985). We have simplified this formulation by assuming that each component has the same scale s_ρ . Furthermore, c_ρ^{-1} is a constant, that when multiplied by s_ρ creates the “spike” component that is narrowly peaked at zero. Note also that by “scale” we are referring to the standard deviation. This formulation effectively allows for sampling from a null model ($k = 1$), a model with a positive constraint ($k = 2$), or a model with a negative constraint ($k = 3$). s_ρ and c_ρ are determined by the researcher. The former can be chosen to reflect a hypothesized effect size, whereas the latter is used to create a null region that is practically equivalent to zero. Examples of this mixture prior are provided in Figure 1 (panel A).

Hypothesis Formulation

The mixture defined in Equation (8) allows for confirmatory hypothesis testing, which is a central contribution of this work. For example, there is often a positive relationship between the mean and variance (Figure 1 panel C). This scientific expectation can be tested as it relates to the mixture components, that is,

$$\begin{aligned}\mathcal{H}_1 : k = 1 \\ \mathcal{H}_2 : k = 2 \\ \mathcal{H}_3 : k = 3.\end{aligned}\tag{9}$$

Recall that $k = 2$ corresponds to the positive constraint. This can then be tested against its compliment \mathcal{H}_c : “not \mathcal{H}_2 .” The corresponding Bayes factor follows,

$$BF_{2c} = \frac{Pr(k = 2 | \mathbf{Y})}{1 - Pr(k = 2 | \mathbf{Y})} / \frac{Pr(k = 2)}{1 - Pr(k = 2)},\tag{10}$$

where the numerator is the posterior odds and the denominator is the prior odds. This expression parallels the informative testing literature (Hojtink, 2011), where hypotheses are often tested against their compliment (e.g., Equation 8 in van de Schoot, Verhoeven, & Hoijtink, 2013). Our mixture approach can also be used for one-sided hypothesis testing, that is,

$$\begin{aligned}\mathcal{H}_0 : k = 1 \\ \mathcal{H}_1 : k = 2,\end{aligned}\tag{11}$$

where $k = 1$ is the “spike” component. In this case, the posterior odds is given as $Pr(k = 2 | \mathbf{Y}) / Pr(k = 1 | \mathbf{Y})$. Because we assumed equal prior odds, this also corresponds to the Bayes factor BF_{10} . Furthermore, our mixture approach can seamlessly be extended to test joint hypotheses that include several correlations. We provide an example of this below. These are major and novel contributions.

At this point, it is important to note that we are focusing on the marginal posterior probabilities of each mixture component, that is, $Pr(k|Y)$, $k = 1, 2, 3$. These are known as posterior inclusion probabilities (PIP). In this example, because we are only testing one parameter, they also correspond to the respective model probabilities. In the case of testing several correlations, however, it is possible to obtain the highest posterior model (HPM). A limitation of this approach is that, with several test relevant parameters, the size of model space becomes prohibitively large. For this reason, we primarily focus on the marginal PIPs (e.g., Table 1 in Wagner & Duller, 2012) and Bayes factors (e.g., p. 216 in Peterson, Swartz, Shete, & Vannucci, 2013).

Software and Estimation

This paper includes a variety of fitted models. To avoid redundancy, we detail their estimation in this section. All models were fitted with the R package **hypMuVar**, which serves as a front-end to the popular Bayesian software JAGS (Plummer, 2013). For demonstrative purposes, we set $s_\rho = 0.5$ and $c_\rho = 50$ (Figure 1, right plot in panel A). Hence, the standard deviation of the “spike” component was 0.01. Each fitted model included four chains of 25,000 iterations each, resulting in a total of 100,000 samples from the posterior distribution. Those samples were saved from an initial adaption phase of 5,000 iterations for each chain. This number of samples provided a good quality of the parameter estimates in which the models converged with potential scale reduction factors \hat{R} smaller than 1.1 (Brooks & Gelman, 1998). The posterior distributions are summarized with means, standard deviations, and 90% credible intervals (CrI). All computations were done in R version 3.6.1 (R Core Team, 2017).

Results

The parameter estimates are displayed in Figure 1. Panel B (left) includes the hierarchical reaction time means for each individual. The pink bars denote 90 % CrIs that excluded the fixed effect or population average (dotted line), which revealed that there were substantial individual differences. Panel B (right) includes the hierarchical reaction time standard deviations (SD) for each individual. This is a unique aspect of the MELSM (Equation 2). Paralleling the mean structure, there were also substantial individual differences in consistently responding. In other words, the SD of the response time distributions varied from individual to individual. Furthermore, as revealed in panel B, there was a 3.4 fold increase from the most to least consistent individuals. On the other hand, there was a 1.8 fold difference in the reaction time means.

In Wagenmakers and Brown (2007), where a lawful relationship between reaction time means and SD s was proposed, a null hypothesis significance test was used in a two-stage approach. This was accomplished by computing the reaction time means and SD s for each person in step-one, and then correlating those estimates in step-two. Our model formulation also provides this correlation (Equation 3), but, importantly, for the hierarchical estimates. The individual means and SD s are displayed in Figure 1 (panel C). “Hierarchical” (grey) refers to the estimates obtained from the MELSM, whereas “Empirical” (orange) refers to the sample based estimates.¹ In these data, the correlation was slightly larger for the hierarchical estimates (0.63 vs. 0.65). The results also highlight the central idea behind

shrinkage, in that the model based estimates are often closer to the fixed effect. This gravitation towards the average was especially pronounced for the larger values.

Panel D includes the posterior distribution for ρ_{01} (Equation 4). Recall that the prior for ρ_{01} is a mixture of three components (e.g., panel A) and the proportion of posterior samples in each can be used to compute Bayes factors. In this case, because $Pr(k = 2|\mathbf{Y}) = 1$, this results in an infinite Bayes factor in favor of a positive correlation. This is decisive evidence (Kass & Raftery, 1995). Hence, the slowest individuals tended to be the least consistent.

Summary

This simple example illustrated several benefits of our innovative approach for characterizing mean–variance relations. For example, the hierarchical formulation reduced variability in the estimates and the correlation was tested with a novel strategy. Of course, individual differences are commonly studied in relation to, say, an experimental effect. In the following section, we thus extend this formulation to accommodate both random intercepts and slopes.

Illustrative Examples

We first return to the Stroop and then proceed to the Flanker inhibition task. In these tasks, there is a congruent and an incongruent condition, and researchers are most interested in the contrast between these conditions. For this contrast defines the effect, and in the data we explore, it defines how well each individual can inhibit irrelevant information. We use illustrative data that were first used in von Bastian et al. (2016), and were then reanalyzed in Haaf and Rouder (2017). The latter included a customary, location-only, individual differences model that included random intercepts and slopes. Note that we are answering a much different question than Haaf and Rouder (2017), which focused on the direction of the random slopes (i.e., the experimental effects). We are explicitly interested in the covariances between the individual difference parameters for the location and scale.

Case 1: The Stroop Task

In this task, participants were asked to count the number of characters displayed with key strikes. This is the *number* Stroop task. For the congruent condition, the number of characters matched the digits displayed (e.g., 3 characters presented as 333). For the incongruent condition, there was a mismatch between the number of characters and the digits displayed (e.g., 2 characters presented as 44). There were 121 participants in total. Each completed 48 trials for the two conditions. We chose these data because there is a large effect, with a mean difference of 65 ms, and also significant individual differences in the “Stroop effect” $\chi^2(3) = 10.94, p = 0.012$.²

¹These were computed as the mean and standard deviation for each person.

²We compared a random intercept model to a random slope model with the R package lme4

Data Set 2: The Flanker Task

In this task, the goal was to identify a vowel (e.g., A or E) or consonant (e.g., B or C). The target was located in the middle, and was "flanked" by two characters on either side. The congruent flankers surrounded the target with letters from the same category (e.g., UUAUU). The incongruent flankers were surrounded by mismatched letters (e.g., CCACC). There was also a neutral condition (##A##), but we only used the congruent and incongruent trials. There were 121 participants in total. Each completed 48 trials for the two conditions. We explicitly chose these data to contrast the Stoop data. The mean difference was small (2 ms) and there were no significant individual differences $\chi^2(3) = 2.03, p = 0.566$. This may seem paradoxical on the "surface," but this demonstrates the utility of extended inference beyond the mean structure.

Model Parameterization

Mean Structure.—We fit the same model to both data sets. The outcome is reaction time for correct responses on the seconds scale, predicted by the experimental condition. For each outcome y , the location sub-model of the response times for the i th person and j th trial is given as

$$y_{ij} \sim \beta_0 + \beta_1(\text{Incongruent}_{ij}) + u_{0i} + u_{1i}(\text{Incongruent}_{ij}) + \epsilon_{ij}. \quad (12)$$

where β_0 is the fixed effect intercept, which is the average reaction time for the congruent condition (the reference category). The reaction times are predicted by the experimental condition (congruent vs. incongruent). Hence, β_1 is the average difference from the congruent condition (i.e., the experimental effect). There are random intercepts u_{0i} that capture the individual deviations from β_0 . For person i , their mean reaction time for the congruent condition is $\beta_0 + u_{0i}$. Additionally, there are random slopes, u_{1i} that capture the individual deviations from β_1 . Hence each person has an experimental effect. Again for person i , their experimental effect is $\beta_1 + u_{1i}$.

Residual Variance Structure.—The above is a traditional individual differences model, in that the "errors" are not modeled. This is not the case for the MELSM, that is,

$$\begin{aligned} \epsilon_{ij} &\sim \mathcal{N}(0, \sigma_{\epsilon_{ij}}^2) \text{ with} \\ \sigma_{\epsilon_{ij}}^2 &= \exp(\eta_0 + \eta_1(\text{Incongruent}_{ij}) \\ &\quad + u_{2i} + u_{3i}(\text{Incongruent}_{ij}))^2. \end{aligned} \quad (13)$$

The subscripts i and j denote residuals for the i th person and j th trial. Furthermore, the residual variance $\sigma_{\epsilon_{ij}}^2$ is a function of a hierarchical model that includes both fixed and random effects. This is a defining feature of the MELSM. The interpretation differs from the random intercepts only model, because there is a predictor in the location sub-model (Equation 12). Hence, this model captures systematic patterns in the residual variance that were not explained by the experimental manipulation. The parameters are analogous to those in Equation (12), but they are on the logarithmic scale. η_0 is the average within-person

variance for the congruent condition. η_1 is the average difference, in residual variance, from the congruent condition. This is the experimental effect on consistently inhibiting irrelevant information. There are also random intercepts u_{2i} and slopes u_{3i} that capture the individual deviations from the fixed effects. For person i , the within-person variance for their congruent responses is $\eta_0 + u_{2i}$. Additionally, $\eta_1 + u_{3i}$ is the experimental effect on within-person variance for person i . This leads to an intuitive interpretation as percentage changes. Note that this formulation essentially generalizes the possibility of heterogeneous residual variance to the individual level.

Random Effects Covariance Matrix.—We assume that the random effects, for both the location and scale sub-models, are drawn from a common multivariate normal distribution

$$\begin{bmatrix} u_{0i} \\ u_{1i} \\ u_{2i} \\ u_{3i} \end{bmatrix} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}), \quad (14)$$

with mean zero and $\mathbf{\Sigma}$ the covariance matrix that includes the random effects variances and covariances. We again use the separation strategy. This can be written as

$$\mathbf{\Sigma} = \text{diag}(\boldsymbol{\tau}) \mathbf{\Omega} \text{diag}(\boldsymbol{\tau}), \quad (15)$$

where $\boldsymbol{\tau}$ is a 4×1 vector, $\tau_j = 0, 1, 2, 3$, that contains the random effects SD s, $\text{diag}(\boldsymbol{\tau}) = \text{diag}(\tau_0, \tau_1, \tau_2, \tau_3)$ is a diagonal matrix with the diagonal elements $\boldsymbol{\tau}$, and $\mathbf{\Omega}$ is a 4×4 correlation matrix.

Standard Deviations.: The SD s capture variability in the random intercepts τ_0 and slopes τ_1 for the mean structure, as well as the random intercepts τ_2 and slopes τ_3 for the residual variance structure. For example, τ_1^2 is the variance in the experimental effects for the location sub-model. This is commonly tested to determine whether there is a “common effect” (Haaf & Rouder, 2017). In this paper we do not test the variance components, because we are explicitly interested in covariance between the random effects across the location and scale sub-models. We return to the topic of testing the random effects variances in the discussion.

Correlations.: The correlation matrix is defined as

$$\mathbf{\Omega} = \begin{bmatrix} 1 & \rho_{01} & \rho_{02} & \rho_{03} \\ \rho_{01} & 1 & \rho_{12} & \rho_{13} \\ \rho_{02} & \rho_{12} & 1 & \rho_{23} \\ \rho_{03} & \rho_{13} & \rho_{23} & 1 \end{bmatrix}. \quad (16)$$

ρ_{01} is the correlation for the random effects *within* the location sub-model. That is, the individual means for the congruent responses and the individual mean differences (i.e., the experimental effect). ρ_{23} is the correlation for the random effects *within* the scale sub-model. This provides the relation between the within-person variances for the congruent responses

and the experimental effects on within-person variance. The remaining correlations that are located in $\Omega_{3:4,1:2}$ capture the relations *across* the location and scale sub-models. Because they are the primary focus in this work, and they lead to novel inferences, we explain them in detail here.

1. ρ_{02} : The correlation between random intercepts. That is, the individual means u_{0i} and within-person variances u_{2i} for the congruent responses. A similar relation is displayed in Figure 1 (panel D), but importantly, in this case, it captures residual variance that can be interpreted as response time consistency.
2. ρ_{03} : The correlation between random intercepts for the location sub-model and random slopes for the scale sub-model. That is, the individual means u_{0i} in the congruent condition and the individual differences u_{3i} in within-person variance, compared to the congruent condition (the effect on response time consistency). Interestingly, this allows for testing whether individuals with the fastest (or slowest) congruent responses tended to have more (or less) consistent incongruent responses (compared to their congruent responses).
3. ρ_{12} : The correlation between random slopes for the location sub-model u_{1i} and random intercepts for the scale sub-model u_{2i} . This is similar to the previous, but the question asked is slightly different. In this case, the correlation captures whether those with the largest (or smallest) effects on their reaction times tended to have more (or less) consistent congruent responses.
4. ρ_{13} : The correlation between the random slopes for the location and scale sub-models. That is, the individual “Stroop/flanker effects” on response time and within-person variance. This is perhaps the most interesting relation for the reaction time modeling literature in particular, because slower individuals are predicted to be more variable (Wagenmakers & Brown, 2007). In this case, it is not whether slower people are less consistent in general, but whether those that struggled to inhibit irrelevant information also inhibited irrelevant information inconsistently.

Prior Specification.—We do not test the fixed effects, and thus use diffuse priors, $\beta_0, \beta_1, \eta_0, \eta_1 \sim \mathcal{N}(0, 5)$. Similarly, we also use diffuse prior distributions for the random effects SDs , $\tau_i \sim \text{Student-}t^+(\mu = 0, \sigma = 1, \nu = 10)$, $i = 0, 1, 2, 3$. These are the same prior used in the random intercepts only model. Note that these are more informative than so-called noninformative prior distributions that customarily set the scale to a large value (e.g., 1000). In more complex models, such as the MELSM, more informative priors can improve sampling efficiency (Gelman, Simpson, & Betancourt, 2017). This is the primary motivation for these priors.

We test the correlations that capture the relations across the location and scale sub-models. Accordingly, a mixture prior distribution is not used for ρ_{01} and ρ_{23}

$$\begin{aligned}
 \rho_{01} &= F^{-1}(z_{01}) \\
 \rho_{23} &= F^{-1}(z_{23}) \\
 z_{01} &\sim \mathcal{N}(0, 0.5) \\
 z_{23} &\sim \mathcal{N}(0, 0.5),
 \end{aligned}
 \tag{17}$$

which include the within sub-model relations. Recall that $F^{-1}(z)$ computes the inverse of the Fisher z transformation for correlations. This offers increased flexibility compared to, say, a uniform distribution. For the remaining correlations, that capture the mean and within-person variance relations, we employ a three component mixture, that is,

$$\begin{aligned}
 \rho_{ij} &\sim F^{-1}(z_{ij}) \\
 z_{ij} &\sim \begin{cases} \mathcal{N}(0, 0.01) & \text{if } k = 1 \\ \mathcal{N}^+(0, 0.50) & \text{if } k = 2 \\ \mathcal{N}^-(0, 0.50) & \text{if } k = 3 \end{cases}
 \end{aligned}
 \tag{18}$$

$$k_{ij} \sim \text{Cat}(\boldsymbol{\pi}), k \in \{1, 2, 3\} \tag{19}$$

where $i = 0, 1$ and $j = 2, 3$ ($i \neq j, i < j$). The standard deviations of the ‘‘spike’’ ($k = 1$) and slab ($k = 2, 3$) components were set to 0.01 and 0.50, respectively. Furthermore, $\boldsymbol{\pi}$ is a 1×3 vector of prior probabilities for each category, such that $\sum_{k=1}^K \pi_k = 1$, and k_{ij} is the indicator for each correlation. We again assume equal prior probabilities $1/K = 0.33$.

These prior distributions were also used in the random intercepts only model. However, there is an important difference. This correlation matrix is larger (4×4 vs. 2×2), which results in restrictions on the size of the correlations. Accordingly, constraints are needed to ensure that the resulting covariance matrix is positive semi-definite. This is automatically handled in JAGS, as invalid values are rejected before the likelihood is evaluated. Hence, the prior in Equation (18) is necessarily restricted to positive semi-definite covariance matrices. In this case, setting the scale to 0.50 ensures that the prior is mostly unaffected by this restriction. Note that Frühwirth-Schnatter and Tüchler (2008) used a Cholesky decomposition to overcome this issue. However, this presents a more serious challenge because the tested elements do not correspond to the covariances (or correlations). This is described in Pinheiro and Bates (see section 2.1, 1996).

Joint Hypotheses

Our mixture approach allows for testing joint hypotheses that include several correlations. For example, a researcher might expect all the correlations to be positive. This can be expressed as

$$\begin{aligned}
 \mathcal{H}_1 &: \text{‘‘}k = 2 \text{ for all correlations’’} \\
 \mathcal{H}_c &: \text{‘‘not } \mathcal{H}_1 \text{’’},
 \end{aligned}
 \tag{20}$$

which compares \mathcal{H}_1 to its complement \mathcal{H}_c . As described above (Equation 10), the Bayes factor BF_{1c} can be computed directly from the posterior samples that are compatible with

\mathcal{H}_1 . This general approach can seamlessly be extended to, say, also considering a null model \mathcal{H}_2 : “ $k = 1$ for all correlations” and redefining the compliment as \mathcal{H}_c : “not \mathcal{H}_1 or \mathcal{H}_2 .” In this example, we test the hypothesis in Equation (20). We also present the marginal posterior probabilities for each correlation.

Results

Comparison to a Mixed-Effects Model.—We first compared the MELSM to a traditional mixed-effects *location* model (MELM). The models are the same (including the priors), but for the MELM, only the mean structure was specified (Equation 12). Further, it is important to note the MELM can be understood as predicting the residual variance with an intercept (η_0). Hence, the remaining scale parameters and covariances within and across the location and scale sub-models are implicitly constrained to zero. For purely descriptive purposes, we use the 90% CrIs to infer whether an individual differs from the average effect.

Figure 2 (panel A) includes the random slopes for the mean structure (i.e., $\beta_1 + u_{1i}$). These correspond to the experimental effects for each person. There are clear differences between the MELM and MELSM. In reference to the empirical estimates (grey line), there was more shrinkage towards β_1 for the MELM. This was non-trivial, in that each model would lead to a different conclusion. On the one hand, because only two individuals differed from the average for the MELM, there does not appear to be notable individual differences in the Stroop task. However, this assumes that the residual variance is a fixed, non-varying constant. On the other hand, with the MELSM, 24 % of the individual effects differed from the average (β_1). These deviations were perhaps small, but they are nonetheless important to consider. These differences between models can be understood in reference to Best Linear Unbiased Prediction (BLUP) of random effects, which are computed assuming a common residual variance (Equation 2.16 in McCulloch, 2003). We emphasize this is not the case for the MELSM.

A similar pattern was revealed for the Flanker task, but the differences between models were even more pronounced. Note that, on average (β_1), there did not appear to be an experimental effect (Table 1). This seemed to generalize to the individual level for the MELM, but not for the MELSM. Indeed, $\approx 10\%$ of the sample had a negative effect that indicates faster incongruent responses, whereas $\approx 15\%$ had an effect in the expected direction (i.e., slower incongruent responses). There were also individual differences in the MELSM, but not the MELM, in that 24% of the sample differed from the fixed effect β_1 .

Hierarchical Shrinkage. We were surprised by the aggressive shrinkage provided by the MELM. We thus took a closer look at the Flanker task. The shrunken estimates for each model are visualized in Figure 2 (panel B). In the MELM, the random slopes in particular were shrunken towards a common value (β_1), whereas, in the MELSM, the estimates were shrunken towards both β_0 and β_1 . Furthermore, due to the aggressive shrinkage, the MELM based estimates could be suggested to reflect mere measurement error (i.e., “noise”). Conversely, by examining the residual variance structure, this can provide key insights into behavioral consistency or stability. As we describe below, there is actually more individual variation in the “noise,” or response time consistency, than in the response times themselves.

Case 1: The Stroop Task.

Mean Structure.: The fixed and random effects are reported in Table 1. The experimental effect on reaction time was 65 ms (Incongruent (β_1) in Table 1). This reproduced the result in Haaf and Rouder (2017). This demonstrates the MELSM and MELM provide similar estimates for the average effect. However, as we described above, there were notable differences at the individual level which is the primary focus of this work.

Residual Variance Structure.: Figure 3 (panel A, “Scale Intercepts”) displays the individual, within-person variances, for the congruent condition (i.e., the random intercepts). There were clear individual differences in response time consistency, which is unique information provided by the MELSM. The individual estimates ranged from 0.06 to 0.31. Expressed on the variance scale, this results in a 25 fold increase from the most to least consistent individuals! To put this in perspective, the minimum–maximum range was 0.49–0.95 (ms) for the response time means. This corresponds to a (nearly) two-fold increase.

Figure 3 (panel A, “Scale Slopes”) displays the individual contrasts, in within-person variance, compared to the congruent condition. These are the experimental effects on response time consistency or, substantively, on consistently inhibiting irrelevant information. The fixed effect (η_1) was 0.16 (Table 1), which indicates a 17 % increase in IIV on average ($[\exp(0.16) - 1] \times 100$). This does not generalize to the individual level. In fact, 24 % of the sample had *more* consistent incongruent responses (i.e., a reduction in residual variance). Further, the change in within-person variance, ranged from a decrease of 28% to an increase of 115%! At the same time, the experimental effects on response time were positive for all subjects (Figure 2 panel A). Hence, while all individuals had slower incongruent responses, some had less variable incongruent responses. This runs contrary to the lawful mean–variance relation seen in Figure 1 (panel A).

Mean–Variance Relations.: The scatter plots in Figure 3 (panel B) displays the individual difference parameters across the mean and residual variance sub-models. Recall that this is the primary target for our proposed testing strategy, which allows for *covariance* selection in the distribution of random effects. Here we report the Bayes factor in favor of the direction of the correlation compared to the compliment. Furthermore, the marginal posterior probabilities for each mixture component are provided in Table 2. Panel ρ_{02} is the relation among the reaction time means and residual within-person variances, or IIV, for the congruent responses. This shows a similar pattern as Figure 1 (panel D), such that the slowest individuals tended to be least consistent ($cor[u_{0j}, u_{2j}] = 0.66$, 90% CrI = [0.56, 0.75], $BF_{1c} = Inf$). Interestingly, as revealed in panel ρ_{01} , there was a negative relation among the experimental effects on IIV and the congruent means. In other words, the fastest individuals in the congruent condition tended to have the largest experimental effects on IIV ($cor[u_{0j}, u_{3j}] = -0.37$, 90% CrI = [-0.52, -0.21], $BF_{1c} = 844$). The correlation displayed in panel ρ_{13} is perhaps the most interesting, as it captures the relation among the experimental effects on reaction time and reaction time consistency (i.e., IIV). This effect was the largest ($cor[u_{1j}, u_{3j}] = 0.74$, 90% CrI = [0.59, 0.87], $BF_{1c} = Inf$) and the data are infinitely more likely under the positive model ($k = 2$) than the compliment. Hence, those that were the slowest

at inhibiting irrelevant information tended to be the least consistent at inhibiting irrelevant information.

Case 2: The Flanker Task.

Mean Structure.: The fixed and random effects are reported in Table 1. The experimental effect was practically zero (≈ 2 ms) and the 90 % CrI included zero. Because the MELSM also characterizes the within-person variance structure, this provides a unique opportunity to learn from the data, over and above the mean structure.

Residual Variance Structure.: Figure 4 (panel A, “Scale Intercepts”) displays the individual, within-person variances, for the congruent condition (i.e., the random intercepts). There were clear individual differences in response time consistency. The response time IIVs ranged from 0.05 to 0.43. Expressed on the variance scale, this results in a 65 fold increase from the most to the least consistent individual ! On the other hand, the minimum-maximum range was 0.43–0.79 (ms) for the response time means. This corresponds to a (nearly) two-fold increase.

In contrast to the Stroop task, the experimental effect on within-person variance was small and the interval included zero (Table 1; Incongruent (η_1)). Recall that this work is explicitly focused on the individual-specific effects. In this case, the fixed effect did not generalize to a large portion of the sample For example, when examining the CrI in reference to zero, 21 % showed an increase in within-person variance. Substantively, this can be interpreted as there being individual differences in consistently inhibiting irrelevant information.

Mean–Variance Relations.: The scatter plots in Figure 4 (panel B) displays the individual difference parameters across the mean and residual variance sub-models. Furthermore, the marginal posterior probabilities for each mixture component are provided in Table 2. Panel ρ_{02} is the relation among the reaction time means and residual within-person variances, or IIV, for the congruent responses. This shows a similar pattern as Figure 1 (panel D) and the Stroop task, such that the slowest individuals tended to be least consistent ($cor[u_{0j}, u_{2j}] = 0.79$, 90% CrI = [0.72, 0.86], $BF_{1c} = \text{Inf}$). Interestingly, while there was a negative relation among the experimental effects on IIV and the congruent means (panel ρ_{01}), the evidence was not strong ($cor[u_{0j}, u_{3j}] = -0.11$, 90% CrI = [-0.23, 0.02], $BF_{1c} = 1.74$). This highlights how our method can be used to compare competing hypotheses. For example, the Bayes factor in favor of the null hypothesis $\mathcal{H}_0 : k = 1$ compared to, say, a negative correlation $\mathcal{H}_1 : k = 3$, is $BF_{01} = 1.5$. The correlation displayed in panel ρ_{13} captures the relation among the experimental effects on reaction time and reaction time consistency (i.e., IIV). This effect was again the largest ($cor[u_{1j}, u_{3j}] = 0.84$, 90% CrI = [0.69, 0.94], $BF_{1c} = \text{Inf}$), which indicates that the data are infinitely more likely under the positive component ($k = 2$) than the compliment. Hence, at the individual level, those that were the slowest at inhibiting irrelevant information tended to be the least consistent at inhibiting irrelevant information. This reproduces the result in the Stroop task.

Joint Hypothesis Testing

An important aspect of our proposed methodology is confirmatory hypothesis testing. As an example, we tested the hypothesis given in Equation (20) that predicts all of the correlations are positive against its complement. For the Stroop task, it was not possible to compute the Bayes factor. This is because the hypothesized model was not visited once which indicates that the hypothesis was not supported by the data. We return to this in the discussion. For the Flanker task, while the data were more likely under \mathcal{H}_c , the Bayes factor is typically considered inconclusive ($BF_{1c} = 0.56$, Kass & Raftery, 1995)

Highest Probability Model

The posterior model probabilities are presented in Table 3. We computed the Bayes factors for the top four models compared to the HPM (\mathcal{M}_1). Hence, the reported value indicates how much more likely the data are under the HMP. For the Stroop task, there was some uncertainty as to the “best” model. This can be seen from the comparison to \mathcal{M}_2 ($BF_{12} = 1.54$). There was even more uncertainty for the Flanker task. In fact, compared to the top four models, the evidence in favor of \mathcal{M}_1 was largely inconclusive. However, in both tasks, ρ_{01} and ρ_{13} were included in the five most probable models. These are the correlations between the random intercepts (the congruent means and residual variances) and slopes (the experimental effect on the means and residual variances).

In the presence of model selection uncertainty (no clear winner), Bayesian model averaging (BMA) can be used (see for example Gronau et al., 2017). This avoids the well-known pitfalls of selecting one model from a candidate set (Hinne, Gronau, Bergh, & Wagenmakers, 2019; Leeb & Pötscher, 2005). We do not explicitly pursue this here, because the marginal posterior distributions were composed of all visited models (not only \mathcal{M}_1), and thus effectively weighted (i.e., the proportion of samples) by the model probabilities. We refer interested readers to Raftery, Madigan, and Hoeting (1997) and Hoeting, Madigan, Raftery, and Volinsky (1999) for a full discussion of BMA. Further Pang and Gill (2009) describes how the “spike” and slab can be used for simultaneous hypothesis testing and model averaging.

Discussion

We introduced a novel framework for testing mean-variance relations in the random effects distribution of hierarchical models. This was accomplished by including a mixed-effects sub-model to predict the residual, within-person variance, that allows for simultaneously estimating individual difference parameters for both the mean (location) and variance (scale) structures. We proposed a three component mixture prior distribution for the explicit goal of *covariance* selection, that is, Bayesian hypothesis testing of mean–variance relations. In several examples, we demonstrated the utility of characterizing the within-person variance structure and testing hypothesized mean–variance relations in the distribution of random effects.

Future Directions

We focused on the person-specific, random effects, which is motivated by recent calls in psychology to place more emphasis at the individual level (Hamaker, 2012; Molenaar, 2004). To this end, we did not test the fixed effects or the random effects variances. However, the presented methodology can be extended to test all aspects of the location and scale sub-models (e.g., Williams, Martin, & Rast, 2019). This would further Rouder, Morey, Speckman, and Province (2012), where the fixed effects can be tested but it is not possible to model the scale or the random effects correlations (i.e., the parameters in Figures 3 and 4 are implicitly constrained to zero). Consequently, the MELSM provides a foundation from which to develop novel Bayesian methodology that can accommodate scale models and, say, allow for testing mean–variance relations among fixed-effects.

We emphasized that residual variance can provide novel insights relating to the consistency of psychological processes. Alternatively, from a measurement perspective, the residuals are mere “error.” It was apparent that treating the residual variance as a fixed, non-varying constant, across individuals was not warranted for a random intercepts model (Figure 1, panel B). This model is commonly used to compute reliability or intraclass correlation coefficients (Bliese, 2000; Shieh, 2016). Our results point towards moving beyond considering reliability as a fixed quantity. In a MELSM, this possibility was briefly described in Hedeker et al. (Equation 10, 2008). Hence, our work motivates further investigating individually varying reliability with the MELSM.

Our results also have theoretical implications for mean–variance lawfulness in psychology. It is important to not only establish “laws,” but also to identify processes that do not adhere to the “law.” For example, there is substantial interest in processes that do not show the typical mean–variance relation (Pratte, Rouder, Morey, & Feng, 2010; Schwarz & Miller, 2012). In the extant literature, this is commonly investigated by averaging across individuals (e.g., delta plots, De Jong et al., 1994). As we demonstrated, our methodology allows for probing mean–variance lawfulness at the individual level. However, we argue that the “law” should first be clarified for within-person experimental effects. For example, if slower reaction times are predicted to necessarily be more inconsistent, several people in both tasks *violate* the “law” (e.g., Figure 3, panel A). On the other hand, if the slowest individuals (compared to the congruent primes) are predicted to have the largest increases in IIV, then the “law” would be *confirmed* in both tasks (Figure 4, panel A). The MELSM can be used to answer this important theoretical question, for example, by extending Haaf and Rouder (2017) to the scale sub-model.

Limitations

When the variances are of interest, it should be noted that their magnitude is also defined by the mean (Baird, Le, & Lucas, 2006; Eid & Diener, 1999; Kalmijn & Veenhoven, 2005; Rouder et al., 2008). In the present work, this particular correlation $cor(u_{0j}, u_{2j})$ was large for both tasks. This can also be seen in Figure 1 (panel D). We refer interested readers to Mestdagh et al. (2018), where a “mean-corrected” measure of variability was introduced. This is not necessarily a “problem” for the MELSM, especially for the other correlations, but it should be considered when making inferences about mean–variance relations.

The responses were right skewed and we assumed normality. This choice is not without precedent (e.g., Haaf & Rouder, 2017; Rouder, Kumar, & Haaf, 2019). In our view, the research question should determine the modeling strategy (Rousselet & Wilcox, 2019). We were explicitly interested in the mean and variance, and in the case of skewed data, both can be estimated accurately when assuming normality (e.g., Figure 1, panel D). There are several distributions that accommodate skew and encode the mean–variance relation (Table 1 in Wagenmakers & Brown, 2007). However, the mean and variance are then a function of the distributional parameters which presents challenges in more complex models (e.g., de Villemereuil, Schielzeth, Nakagawa, & Morrissey, 2016). This also applies to the log-normal distribution, although the location and scale retain their interpretability as the mean and variance of $\log(y)$. Thus, the log-normal distribution seems ideal for those not comfortable assuming normality.³ Furthermore, because the correlations were so striking, we computed the individual means and *SDs*, as well as contrasts between conditions. The same pattern emerged, in that there were *large* correlations for both tasks which reproduced Figures 3 and 4 (panel B). Hence, our results are not an artifact, as they can be seen right in the data.

Conclusion

In the social-behavioral sciences, it is customary to view homogeneous variance as an assumption to satisfy. The central message of this work is to move beyond that perspective, and to view within-person variance as an opportunity to gain a richer understanding of psychological processes. We introduced a framework to facilitate this transition by targeting mean–variance relations. The proposed methodology is not only suited for inhibition tasks, as presented in this work, but can also be used more generally in repeated measures designs. By focusing on the within-person variance, that is response time consistency, this opened up possibilities for modeling a component that is often disregarded as mere “noise.” The illustrative examples highlighted such possibilities and demonstrated that the residual variance may show systematic patterns, individual differences, and relations with the mean structure.

Funded by:

National Institutes of Health, National Institute on Aging

National Science Foundation

References

- Aarts E, Dolan CV, Verhage M, & van der Sluis S (2015). Multilevel analysis quantifies variation in the experimental effect while optimizing power and preventing false positives. *BMC Neurosci*, 16(1), 94. doi: 10.1186/s12868-015-0228-5 [PubMed: 26685825]
- Aarts E, Verhage M, Veenliet JV, Dolan CV, & van der Sluis S (2014, 4). A solution to dependency: using multilevel analysis to accommodate nested data. *Nature Neuroscience*, 17(4), 491–496. doi: 10.1038/nn.3648 [PubMed: 24671065]

³We also fitted log-normal models. The results were very similar to assuming normality.

- Alvarez I, Niemi J, & Simpson M (2014). Bayesian inference for a covariance matrix. In Annual conference on applied statistics in agriculture (Vol. 20, pp. 1669–1696). doi: 10.1214/aos/1176348885
- Baird BM, Le K, & Lucas RE (2006). On the nature of intraindividual personality variability: Reliability, validity, and associations with well-being. *Journal of personality and social psychology*, 90(3), 512. [PubMed: 16594835]
- Barnard J, McCulloch R, & Meng X-L (2000). Modelling Covariance Matrices in Terms of Standard Deviations and Correlations With Applications to Shrinkage. *Statistica Sinica*, 10(4), 1281–1311. doi: 10.2307/24306780
- Barr DJ, Levy R, Scheepers C, & Tily HJ (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255–278. doi: 10.1016/j.jml.2012.11.001
- Bauer DJ (2011,4). Evaluating Individual Differences in Psychological Processes:. *Current Directions in Psychological Science*, 20(2). doi: 10.1177/0963721411402670
- Bliese PD (2000). Within-group agreement, non-independence, and reliability: Implications for data aggregation and analysis (1st ed.; Klein KJ & Kozlowski SW, Eds.). San Francisco: Jossey-Bass.
- Brooks SPB, & Gelman AG (1998, 2). General methods for monitoring convergence of iterative simulations. *Journal of computational and graphical statistics*, 7(4), 434–455. doi: 10.2307/1390675
- Cattell RB, Cattell AKS, & Rhymer RM (1947). P-technique demonstrated in determining psychophysiological source traits in a normal individual. *Psychometrika*, 12(4), 267–288. [PubMed: 18921433]
- Christ BU, Combrinck MI, & Thomas KGF (2018). Both Reaction Time and Accuracy Measures of Intraindividual Variability Predict Cognitive Performance in Alzheimer’s Disease. *Frontiers in human neuroscience*, 12, 124. doi: 10.3389/fnhum.2018.00124 [PubMed: 29686610]
- Cleveland WS, Denby L, & Liu C (2003). Random scale effects. (2), 33. Retrieved from stat.bell-labs.com
- Daniels MJ, Kass RE (1999, 12). Nonconjugate Bayesian Estimation of Covariance Matrices and its Use in Hierarchical Models. *Journal of the American Statistical Association*, 94(448), 1254–1263. doi: 10.1080/01621459.1999.10473878
- Davidson DJ, Zacks RT, & Williams CC (2003, 6). Stroop Interference, Practice, and Aging. *Aging, Neuropsychology, and Cognition*, 10(2), 85–98. doi: 10.1076/anec.10.2.85.14463
- De Jong R, Liang C-C, & Lauber E (1994). Conditional and unconditional automaticity: A dual-process model of effects of spatial stimulus-response correspondence. *Journal of Experimental Psychology: Human Perception and Performance*, 20(4), 731–750. doi: 10.1037/0096-1523.20.4.731 [PubMed: 8083631]
- de Villemereuil P, Schielzeth H, Nakagawa S, & Morrissey M (2016, 11). General methods for evolutionary quantitative genetic inference from generalized mixed models. *Genetics*, 204(3), 1281–1294. doi: 10.1534/genetics.115.186536 [PubMed: 27591750]
- Efron B, & Morris C (1977). Stein’s paradox in statistics. *Scientific American*, 236(5), 119–127.
- Eid M, & Diener E (1999). Intraindividual variability in affect: Reliability, validity, and personality correlates. *Journal of Personality and Social Psychology*, 76(4), 662.
- Fagot D, Mella N, Borella E, Ghisletta P, Lecerf T, & De Ribaupierre A (2018). Intra-Individual Variability from a Lifespan Perspective: A Comparison of Latency and Accuracy Measures. *Journal of Intelligence*, 6(1), 16.
- Fiske DW, & Rice L (1955). Intra-individual response variability. *Psychological Bulletin*, 52(3), 217. [PubMed: 14371891]
- Foulley J, & Quaas R (1995). Heterogeneous variances in Gaussian linear mixed models. *Genetics Selection Evolution*, 27(3), 211. doi: 10.1186/1297-9686-27-3-211
- Foulley JL, San Cristobal M, Gianola D, & Im S (1992). Marginal likelihood and Bayesian approaches to the analysis of heterogeneous residual variances in mixed linear Gaussian models. *Computational Statistics and Data Analysis*, 13(3), 291–305. doi: 10.1016/0167-9473(92)90137-5
- Frühwirth-Schnatter S, & Tüchler R (2008, 3). Bayesian parsimonious covariance estimation for hierarchical linear mixed models. *Statistics and Computing*, 18(1), 1–13. doi: 10.1007/s11222-007-9030-2

- Gelman A (2006). Prior distribution for variance parameters in hierarchical models. *Bayesian Analysis*, 1(3), 515–533. doi: 10.1214/06-BA117A
- Gelman A, Simpson D, & Betancourt M (2017). The prior can generally only be understood in the context of the likelihood. ArXiv preprint.
- George EI, & McCulloch RE (1993). Variable selection via Gibbs sampling. *Journal of the American Statistical Association*, 88(August 2015), 881–889. doi: 10.1080/01621459.1993.10476353
- George EI, & McCulloch RE (1997). Approaches for Bayesian Variable Selection. *Statistica Sinica*, 7, 339–373. doi: 10.1.1.211.4871
- Gronau QF, Van Erp S, Heck DW, Cesario J, Jonas KJ, & Wagenmakers E-J (2017). A Bayesian model-averaged meta-analysis of the power pose effect with informed and default priors: the case of felt power. *Comprehensive Results in Social Psychology*, 2(1), 123–138. doi: 10.1080/23743603.2017.1326760
- Gutiérrez-Peña E, Smith AFM, Bernardo JM, Consonni G, Veronese P, George EI, ... Morris CN (1997, 6). Exponential and bayesian conjugate families: Review and extensions. *Test*, 6(1), 1–90. doi: 10.1007/BF02564426
- Haaf JM, & Rouder JN (2017). Developing Constraint in Bayesian Mixed Models. *Psychological Methods*, 22(4), 779–798. doi: 10.1037/met0000156 [PubMed: 29265850]
- Hamaker EL (2012). Why researchers should think “within-person”: A paradigmatic rationale. In Mehl MR & Conner TS (Eds.), *Handbook of research methods for studying daily life* (1st ed., pp. 43–61). New York.
- Heck DW, Overstall AM, Gronau QF, & Wagenmakers EJ (2018). Quantifying uncertainty in transdimensional Markov chain Monte Carlo using discrete Markov models. *Statistics and Computing*, 29(4), 631–643. doi: 10.1007/s11222-018-9828-0
- Hedeker D, Mermelstein RJ, & 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 [PubMed: 17970819]
- Hedeker D, Mermelstein RJ, & 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 [PubMed: 22419604]
- Hinne M, Gronau QF, Bergh D. v. d., & Wagenmakers E-J (2019). A conceptual introduction to Bayesian model averaging. PsyArXiv. doi: 10.31234/OSF.IO/PFYNB
- Hoeting JA, Madigan D, Raftery AE, & Volinsky CT (1999). Bayesian Model Averaging: A Tutorial. *Statistical Science*, 14(4), 382–417. doi: 10.2307/2676803
- Hoijtink H (2011). *Informative hypotheses: Theory and practice for behavioral and social scientists*. Chapman and Hall/CRC.
- Horn JL (1972). State, trait and change dimensions of intelligence. *British Journal of Educational Psychology*, 42(2), 159–185.
- Hsu CW, Sinay MS, & Hsu JS (2012, 4). Bayesian estimation of a covariance matrix with flexible prior specification. *Annals of the Institute of Statistical Mathematics*, 64(2), 319–342. doi: 10.1007/s10463-010-0314-5
- Huang A, & Wand MP (2013). Simple marginally noninformative prior distributions for covariance matrices. *Bayesian Analysis*, 8(2), 439–452. doi: 10.1214/13-BA815
- James W, & Stein C (1961). Estimation with quadratic loss. *Breakthroughs in Statistics: Foundations and basic theory*(30), 361–379. doi: 10.1177/0278364907080252
- Judd CM, Westfall J, & Kenny DA (2012). Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. *Journal of Personality and Social Psychology*, 103(1), 54–69. doi: 10.1037/a0028347 [PubMed: 22612667]
- Kalin AM, Plfuger M, Gietl AF, Riese F, Jäncke L, Nitsch RM, & Hock C (2014, 7). Intraindividual variability across cognitive tasks as a potential marker for prodromal Alzheimers disease. *Frontiers in Aging Neuroscience*, 6, 147. doi: 10.3389/fnagi.2014.00147 [PubMed: 25071556]
- Kalmijn W, & Veenhoven R (2005). Measuring inequality of happiness in nations: In search for proper statistics. *Journal of Happiness Studies*, 6(4), 357–396.
- Kass RE, & Raftery AE (1995). Bayes Factors. *Journal of the American Statistical Association*, 90(430), 773–795.

- Kuo L, & Mallick B (1998). Variable Selection for Regression Models. *Sankhy : The Indian Journal of Statistics, Series B*, 60(1), 65–81.
- 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.
- Leeb H, & Pötscher BM (2005, 2). Model selection and inference: Facts and fiction. *Econometric Theory*, 21(1), 21–59. doi: 10.1017/S0266466605050036
- Leonard T, & Hsu JSJ (1992). Bayesian Inference for a Covariance Matrix Tom Leonard; John S. J. Hsu. *Annals of Statistics*, 20(4), 1669–1696.
- Leppink J (2019, 3). When Negative Turns Positive and Vice Versa: The Case of Repeated Measurements. *Health Professions Education*, 5(1), 76–81. doi: 10.1016/J.HPE.2017.03.004
- Leppink J, & Merriënboer JGGV (2015). The Beast of Aggregating Cognitive Load Measures in Technology-Based Learning. *Educational Technology & Society*, 18(4), 230–245.
- Lin SP, & Perlman DM (1985). A Monte Carlo comparison of four estimators of a covariance matrix (cannot found online source). *Journal of Multivariate Analysis*.
- Liu S, Rovine MJ, & Molenaar PCM (2012). Selecting a linear mixed model for longitudinal data: Repeated measures analysis of variance, covariance pattern model, and growth curve approaches. *Psychological Methods*, 17(1), 15–30. doi: 10.1037/a0026971 [PubMed: 22251268]
- Luce RD (1986). *Response times*. New York: Oxford University Press.
- MacDonald SWS, Hultsch DF, & Dixon RA (2008). Predicting impending death: inconsistency in speed is a selective and early marker. *Psychology and aging*, 23(3), 595. [PubMed: 18808249]
- Mazur JE (2006, 3). Mathematical models and the experimental analysis of behavior. *Journal of the experimental analysis of behavior*, 85(2), 275–91. doi: 10.1901/JEAB.2006.65-05 [PubMed: 16673829]
- McCulloch CE (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.
- Mestdagh M, Pe M, Pestman W, Verdonck S, Kuppens P, & Tuerlinckx F (2018, 12). Sidelineing the mean: The relative variability index as a generic mean-corrected variability measure for bounded variables. *Psychological Methods*, 23(4), 690–707. doi: 10.1037/met0000153 [PubMed: 29648843]
- Mitchell DJ, Fanson BG, Beckmann C, & Biro PA (2016, 10). Towards powerful experimental and statistical approaches to study intraindividual variability in labile traits. *Royal Society Open Science*, 3(10). doi: 10.1098/rsos.160352
- Mitchell TJ, & Beauchamp JJ (1988). Bayesian variable selection in linear regression (with discussion). *J. Amer. Statist. Assoc*, 83(1988), 1023–1036.
- Molenaar PCM (2004). A Manifesto on Psychology as Idiographic Science: Bringing the Person Back Into Scientific Psychology, This Time Forever. *Measurement: Interdisciplinary Research & Perspective*. doi: 10.1207/s15366359mea0204{_}1
- O’Hara RB, & Sillanpää MJ (2009). A review of bayesian variable selection methods: What, how and which. *Bayesian Analysis*, 4(1), 85–118. doi: 10.1214/09-BA403
- Pang X, & Gill J (2009). Spike and slab prior distributions for simultaneous Bayesian hypothesis testing, model selection, and prediction, of nonlinear outcomes. Unpublished Manuscript.
- Peterson C, Swartz M, Shete S, & Vannucci M (2013, 6). Bayesian Model Averaging for Genetic Association Studies. In Do K, Qin ZS, & Vannucci M (Eds.), *Advances in statistical bioinformatics* (pp. 208–223). New York: Cambridge University Press. doi: 10.1017/cbo9781139226448.011
- Pinheiro JC, & Bates DM (1996). Unconstrained parametrizations for variance-covariance matrices. *Statistics and Computing*, 6(3), 289–296. doi: 10.1007/BF00140873
- Plummer M (2013). JAGS: Just Another Gibbs Sampler. Retrieved from <http://mcmc-jags.sourceforge.net/>
- Pratte MS, Rouder JN, Morey RD, & Feng C (2010, 10). Exploring the differences in distributional properties between Stroop and Simon effects using delta plots. *Attention, Perception & Psychophysics*, 72(7), 2013–2025. doi: 10.3758/APP.72.7.2013

- R Core Team. (2017). R: A Language and Environment for Statistical Computing. Vienna, Austria: R Foundation for Statistical Computing.
- Raftery AE, Madigan D, & Hoeting JA (1997). Bayesian Model Averaging for Linear Regression Models. *Journal of the American Statistical Association*, 92(437), 179–191. doi: 10.1080/01621459.1997.10473615
- Ram N, & Gerstorff D (2009). Time-structured and net intraindividual variability: Tools for examining the development of dynamic characteristics and processes. *Psychology and aging*, 24(4), 778. [PubMed: 20025395]
- 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 SM, & Sparks C (2012a). 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 [PubMed: 26734847]
- Rast P, Hofer SM, & Sparks C (2012b). 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 [PubMed: 26734847]
- Ridderinkhof KR, Scheres A, Oosterlaan J, & Sergeant JA (2005). Delta Plots in the Study of Individual Differences: New Tools Reveal Response Inhibition Deficits in AD/HD That Are Eliminated by Methylphenidate Treatment. *Journal of Abnormal Psychology*, 114(2), 197–215. doi: 10.1037/0021-843X.114.2.197 [PubMed: 15869351]
- Röcke C, & Brose A (2013). Intraindividual Variability and Stability of Affect and Well-Being. *GeroPsych*, 26(3), 185–199. doi: 10.1024/1662-9647/a000094
- Rouder JN, Kumar A, & Haaf JM (2019). Why most studies of individual differences with inhibition tasks are bound to fail. *PsyArXiv*, 1–37.
- Rouder JN, Morey RD, Speckman PL, & Province JM (2012, 10). Default Bayes factors for ANOVA designs. *Journal of Mathematical Psychology*, 56(5), 356–374. doi: 10.1016/j.jmp.2012.08.001
- Rouder JN, Tuerlinckx F, Speckman P, Lu J, & Gomez P (2008). A hierarchical approach for fitting curves to response time measurements. *Psychonomic Bulletin and Review*, 15(6), 1201–1208. doi: 10.3758/PBR.15.6.1201 [PubMed: 19001591]
- Rouder JN, Yue Y, Speckman PL, Pratte MS, & Province JM (2010). Gradual growth versus shape invariance in perceptual decision making. *Psychological Review*, 117(4), 1267–1274. doi: 10.1037/a0020793 [PubMed: 21038978]
- Roussellet GA, & Wilcox RR (2019). Reaction times and other skewed distributions: problems with the mean and the median. *bioRxiv*, 383935. doi: 10.1101/383935
- Schwarz W, & Miller J (2012, 8). Response time models of delta plots with negative-going slopes. *Psychonomic Bulletin & Review*, 19(4), 555–574. doi: 10.3758/s13423-012-0254-6 [PubMed: 22610358]
- 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 [PubMed: 26182855]
- Townsend JT (2008, 10). *Mathematical Psychology: Prospects For The 21 Century: A Guest Editorial*. *Journal of mathematical psychology*, 52(5), 269–280. doi: 10.1016/j.jmp.2008.05.001 [PubMed: 19802342]
- van de Schoot R, Verhoeven M, & Hoijtink H (2013, 1). Bayesian evaluation of informative hypotheses in SEM using Mplus: A black bear story. *European Journal of Developmental Psychology*, 10(1), 81–98. doi: 10.1080/17405629.2012.732719
- von Bastian CC, Souza AS, & Gade M (2016, 2). No evidence for bilingual cognitive advantages: A test of four hypotheses. *Journal of experimental psychology. General*, 145(2), 246–258. doi: 10.1037/xge0000120 [PubMed: 26523426]
- Wagenmakers EJ, & Brown S (2007). On the Linear Relation Between the Mean and the Standard Deviation of a Response Time Distribution. *Psychological Review*, 114(3), 830–841. doi: 10.1037/0033-295X.114.3.830 [PubMed: 17638508]

- Author Manuscript
- Author Manuscript
- Author Manuscript
- Author Manuscript
- Wagner H, & Duller C (2012). Bayesian model selection for logistic regression models with random intercept. *Computational Statistics and Data Analysis*, 56(5), 1256–1274. doi: 10.1016/j.csda.2011.06.033
- Wang H (2015). Scaling it up: Stochastic search structure learning in graphical models. *Bayesian Analysis*, 10(2), 351–377. doi: 10.1214/14-BA916
- Watts A, Walters RW, Hoffman L, & Templin J (2016). Intraindividual 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 DR, Carlsson R, & Bürkner P-C (2017). Between-litter variation in developmental studies of hormones and behavior: Inflated false positives and diminished power. *Frontiers in Neuroendocrinology*(August), 0–1. doi: 10.1016/j.yfrne.2017.08.003
- Williams DR, Martin SR, & 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 DR, & Rast P (2018). A Bayesian Nonlinear Mixed-Effects Location Scale Model for Learning., 1–18.
- Wolsiefer K, Westfall J, & Judd CM (2017, 8). Modeling stimulus variation in three common implicit attitude tasks. *Behavior Research Methods*, 49(4), 1193–1209. doi: 10.3758/s13428-016-0779-0 [PubMed: 27519882]
- Woodrow H (1932). Quotidian variability. *Psychological Review*, 39(3), 245.
- Wright BC (2017, 8). What Stroop tasks can tell us about selective attention from childhood to adulthood. *British journal of psychology* (London, England : 1953), 108(3), 583–607. doi: 10.1111/bjop.12230

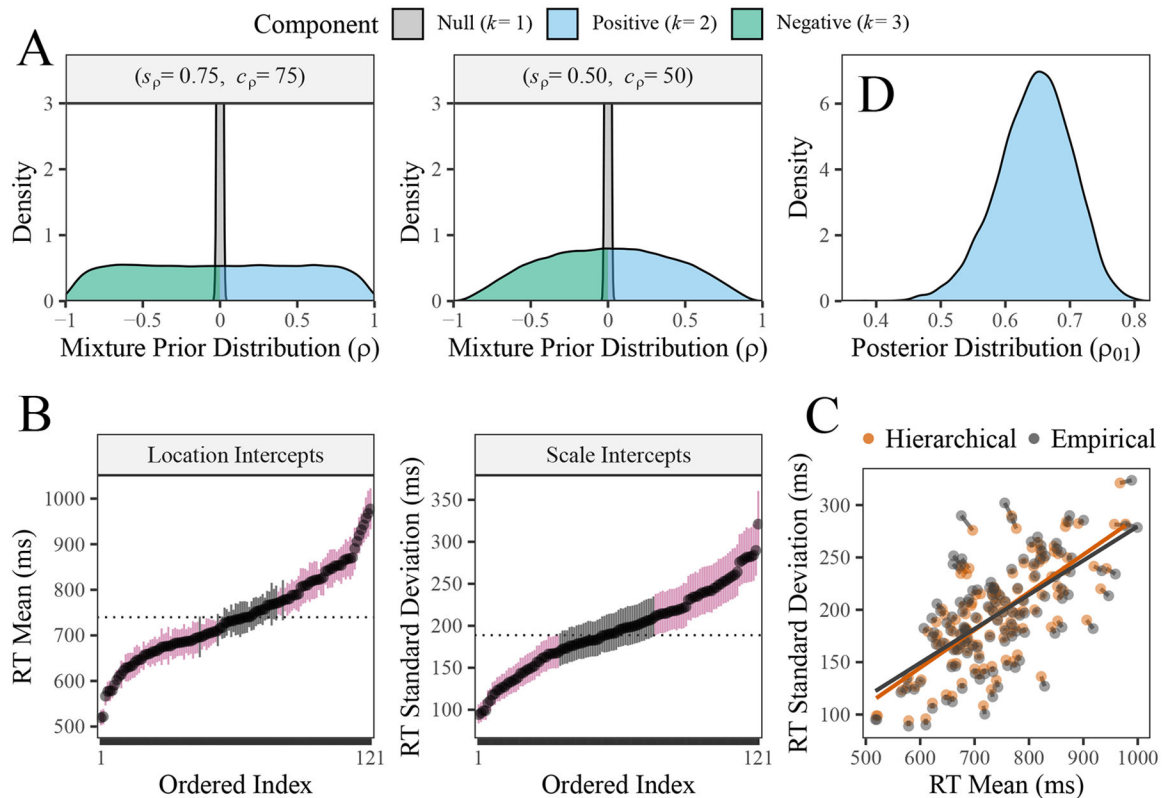
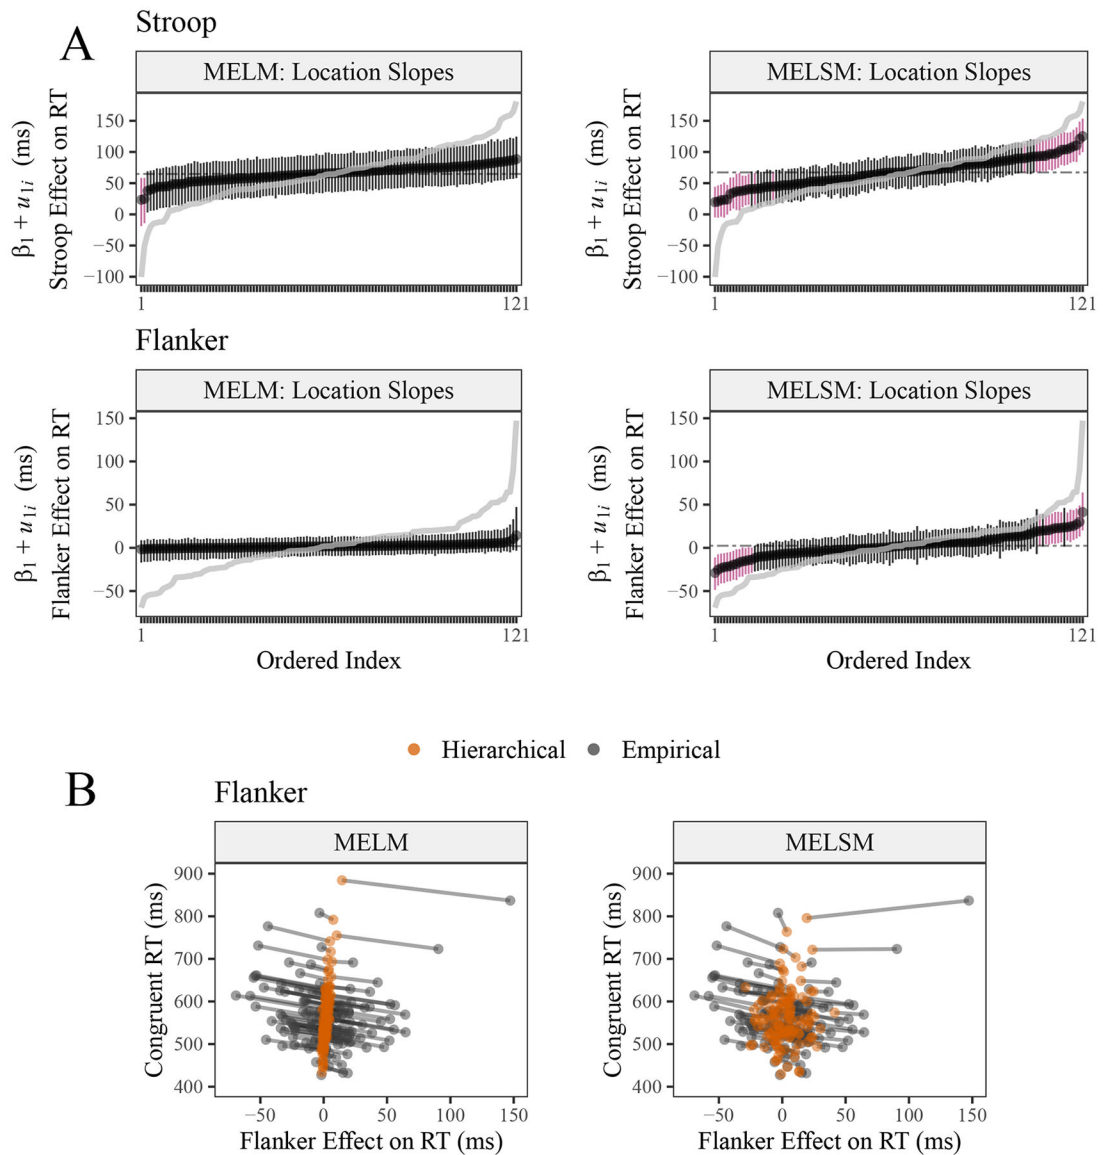
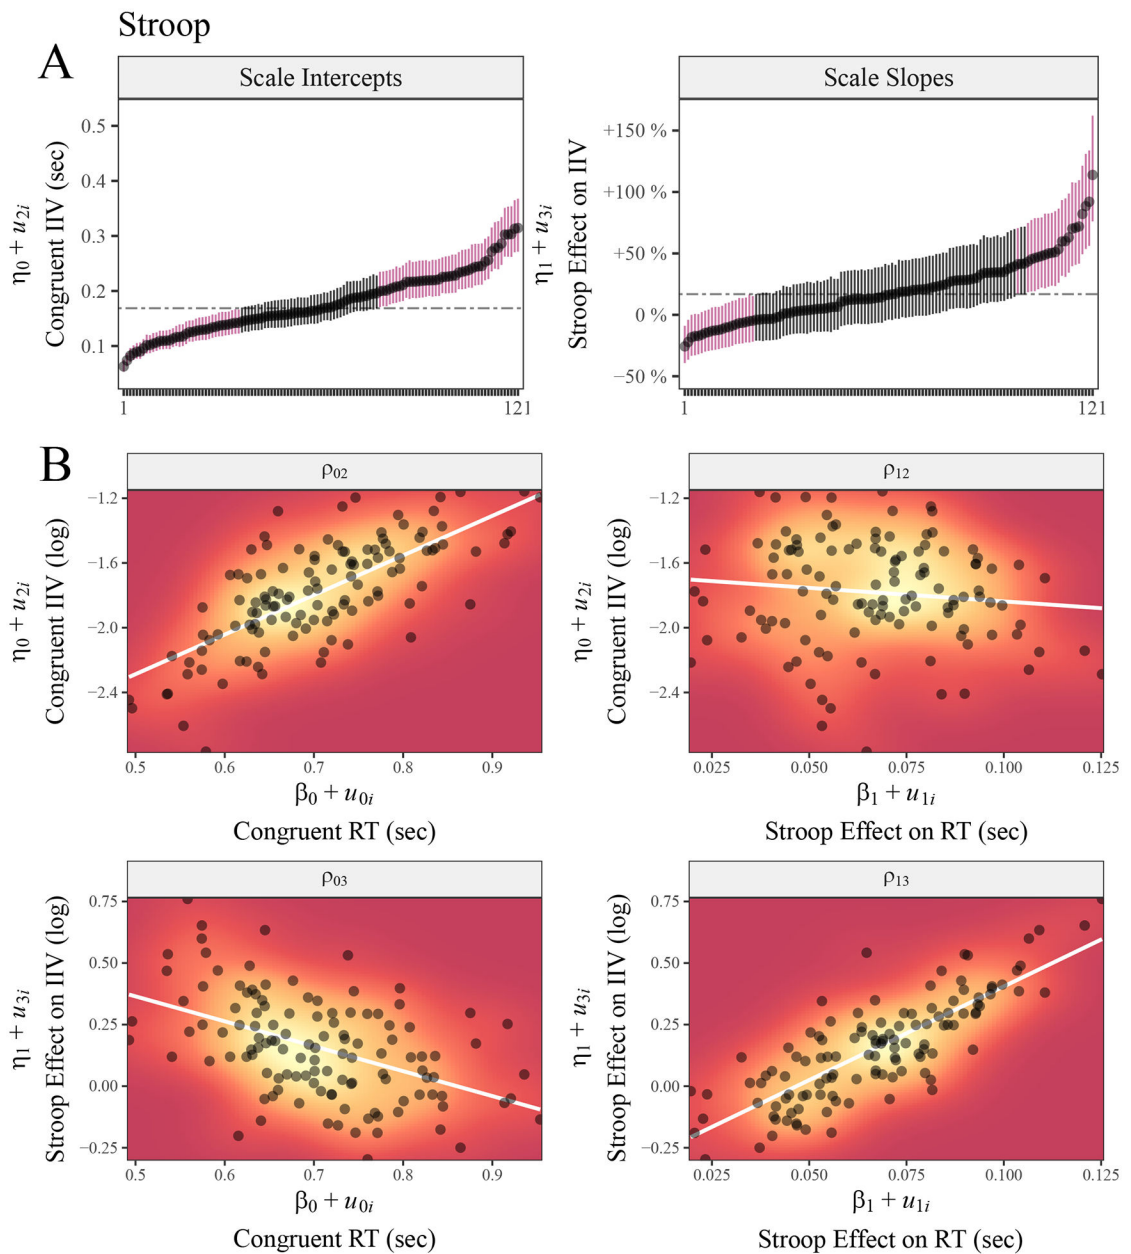


Figure 1.

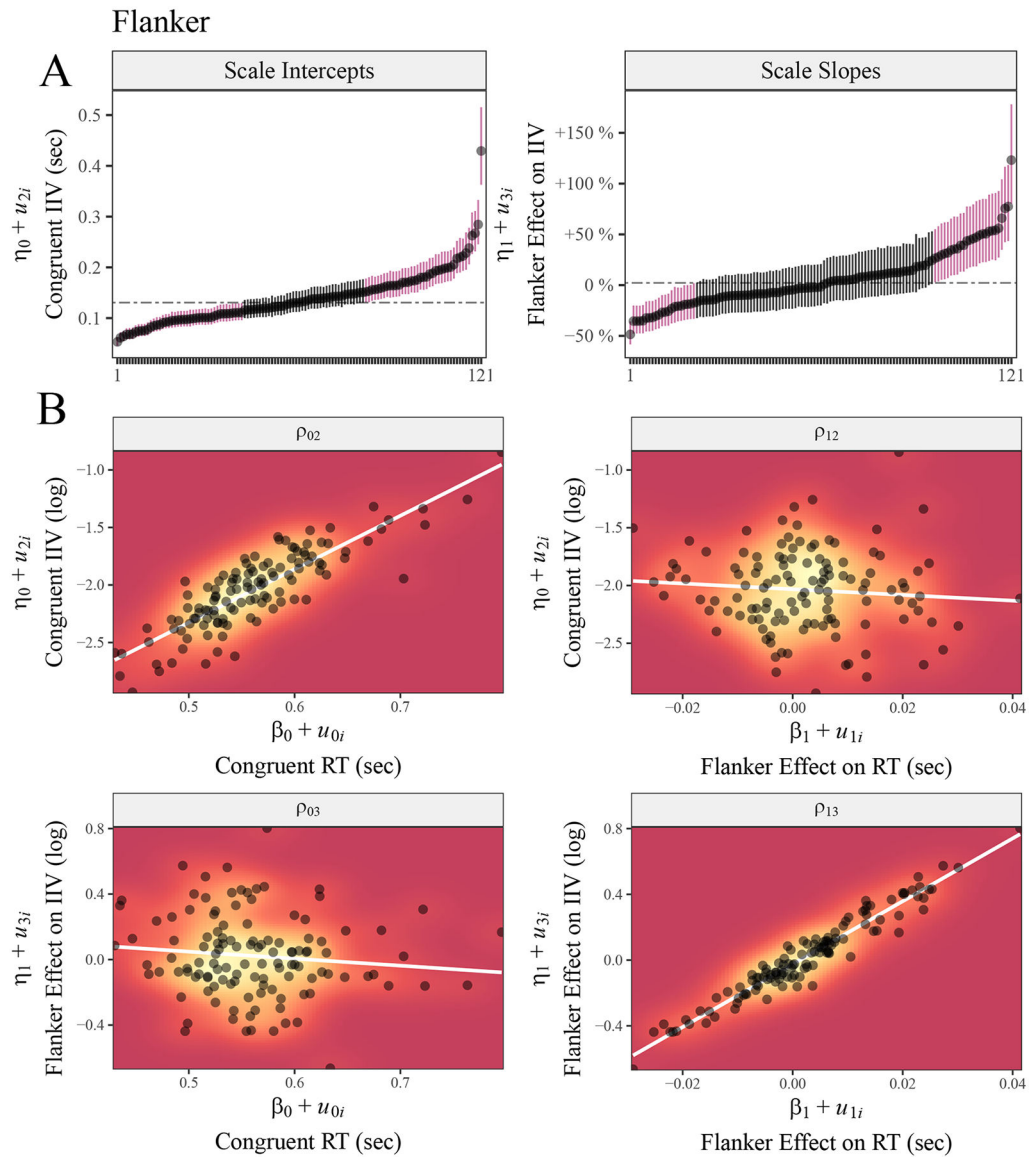
A) Illustrative mixture prior distributions for testing correlations in the distribution of random effects. s_ρ is the scale parameter for the slab ($k = 2, 3$), whereas c_ρ^{-1} is a constant, that when multiplied by s_ρ produces the “spike” ($k = 1$). The basic idea is that the mixture components are competing hypotheses. Hence it is possible to compare, say, $\mathcal{H}_0 : k = 1$ vs. $\mathcal{H}_1 : k = 2$, which results in a one-sided hypothesis test. B) Location and scale intercepts for each individual (black dots). The former is reaction time means and the latter is reaction time standard deviations (SD). The pink error bars denote 90 % credible intervals that excluded the fixed effect (dotted line). This is a key aspect of the MELSM, as location only mixed-effects models customarily assume a common reaction time SD for all subjects. C) Hierarchical shrinkage that effectively smooths the individual estimates towards the fixed effect. The “Hierarchical” estimates (orange) were obtained from the MELSM, whereas the “Empirical” estimates (grey) were obtained from the data by computing the means and SD s for each individual. This reveals that there is hierarchical shrinkage for both the random scale and location effects. D) The posterior distribution for ρ_{01} (Equation 6) that captures the correlation between the individual means and SD s in the distribution of random effects (Equation 3). The blue corresponds to the positive slab ($k = 2$), where $Pr(k = 2 | \mathbf{Y}) = 1.0$, which results in decisive evidence. That is, the slowest individuals tended to be the least consistent.

**Figure 2.**

A) Location random slopes that correspond to the experimental effects on reaction time. MELM is the mixed-effects location only model and MELSM is the mixed-effects location scale model. The pink error bars denote 90 % credible intervals that excluded the fixed effect (dotted line) and the grey lines denote the empirical mean differences between conditions for each person. This plot reveals that the MELM estimates exhibit more shrinkage than the MELSM estimates. This results in less pronounced individual differences, in that, for the MELM, fewer slopes excluded the fixed effect. B) Hierarchical shrinkage for the Flanker task. The stronger shrinkage was especially severe for the experimental effects, as they gravitated towards a common slope in the MELM. The MELSM also provided shrinkage, but it was less pronounced than the MELM.

**Figure 3.**

A) Scale random intercepts and slopes for the Stroop task. The former is the residual within-person variance, or intraindividual variability (IIV), in each subjects' congruent responses. The latter is the difference in IIV between the congruent and incongruent responses. That is, the experimental effect on within-person variance. Exponentiation allows for interpreting the estimates in terms of percentage change in IIV (e.g., +150 % denotes a 150 % increase in IIV compared to the congruent responses). The pink error bars denote 90 % credible intervals that excluded the fixed effect (dotted line). B) Scatter plots of the random effects. Each dot corresponds to a study subject. The correlations correspond to those in Equation (16). The heatmap is a two-dimensional density estimate, where the gradient goes from dark red (low density) to light yellow (high density).

**Figure 4.**

A) Scale random intercepts and slopes for the Flanker task. The former is the residual within-person variance, or intraindividual variability (IIV), in each subjects' congruent responses. The latter is the difference in IIV between the congruent and incongruent responses. That is, the experimental effect on within-person variance. Exponentiation allows for interpreting the estimates in terms of percentage change in IIV (e.g., +150 % denotes a 150 % increase in IIV compared to the congruent responses). The pink error bars denote 90 % credible intervals that excluded the fixed effect (dotted line). B) Scatter plots of the random effects. Each dot corresponds to a study subject. The correlations correspond to those in Equation (16). The heatmap is a two-dimensional density estimate, where the gradient goes from dark red (low density) to light yellow (high density).

Table 1

Parameter estimates

Parameter	Location Model					
	Stroop			Flanker		
	M	SD	90 % CrI	M	SD	90 % CrI
Intercept (β_0)	0.71	0.01	[0.69,0.72]	0.56	0.01	[0.55,0.57]
Incongruent (β_1)	0.07	0.00	[0.06,0.07]	0.00	0.00	[-0.00,0.01]
Random Intercept $SD(\tau_0)$	0.10	0.01	[0.09,0.11]	0.07	0.01	[0.06,0.08]
Random Slope $SD(\tau_1)$	0.03	0.00	[0.02,0.03]	0.02	0.00	[0.01,0.02]

Parameter	Scale Model					
	Stroop			Flanker		
	M	SD	90 % CrI	M	SD	90 % CrI
Intercept (η_0)	-1.78	0.03	[-1.84,-1.73]	-2.04	0.03	[-2.09,-1.98]
Incongruent (η_1)	0.16	0.03	[0.11,0.20]	0.02	0.03	[-0.03,0.07]
Random Intercept $SD(\tau_2)$	0.34	0.02	[0.30,0.39]	0.36	0.02	[0.32,0.40]
Random Slope $SD(\tau_3)$	0.23	0.02	[0.20,0.27]	0.28	0.02	[0.24,0.32]

Note. Posterior mean (M) and standard deviation (SD). The scale model estimates are on the log scale. β_1 is the average experimental effect on response time. η_1 is the average experimental effect on response time consistency.

Table 2

Exhaustive hypothesis testing results

Parameter	Stroop				
	M	SD	$Pr(\mathcal{H}_0 \mathbf{Y})$	$Pr(\mathcal{H}_1 \mathbf{Y})$	$p(\mathcal{H}_2 \mathbf{Y})$
$cor(u_{0j}, t_{2j}) (\rho_{02})$	0.66	0.06	0.00	1.00	0.00
$cor(u_{0j}, t_{3j}) (\rho_{03})$	-0.37	0.10	0.01	0.00	0.99
$cor(u_{1j}, t_{2j}) (\rho_{12})$	-0.04	0.10	0.56	0.11	0.33
$cor(u_{1j}, t_{3j}) (\rho_{13})$	0.75	0.08	0.00	1.00	0.00
Flanker					
Parameter	M	SD	$p(\mathcal{H}_0 \mathbf{Y})$	$p(\mathcal{H}_1 \mathbf{Y})$	$p(\mathcal{H}_2 \mathbf{Y})$
$cor(u_{0j}, t_{2j}) (\rho_{02})$	0.80	0.04	0.00	1.00	0.00
$cor(u_{0j}, t_{3j}) (\rho_{03})$	-0.05	0.09	0.56	0.06	0.38
$cor(u_{1j}, t_{2j}) (\rho_{12})$	-0.08	0.15	0.50	0.10	0.40
$cor(u_{1j}, t_{3j}) (\rho_{13})$	0.84	0.08	0.00	1.00	0.00

Note. Posterior mean (M) and standard deviation (SD). $\mathcal{H}_0: \rho_{ij} \approx 0$ ($k = 1$) vs. $\mathcal{H}_1: \rho_{ij} > 0$ ($k = 2$) vs. $\mathcal{H}_2: \rho_{ij} < 0$ ($k = 3$). Random location μ_{0j} and scale σ_{1j} intercepts. Random location μ_{2j} and scale σ_{3j} slopes.

Table 3

Top five most probable models

Model	Stroop					Flanker					
	BF_{1i}	ρ_{02}	ρ_{03}	ρ_{12}	ρ_{13}	Model	BF_{1i}	ρ_{01}	ρ_{02}	ρ_{12}	ρ_{13}
$\mathcal{M}_1^{\text{HPM}}$	–	$k=2$	$k=3$	$k=1$	$k=2$	$\mathcal{M}_1^{\text{HPM}}$	–	$k=2$	$k=1$	$k=1$	$k=2$
\mathcal{M}_2	1.79	$k=2$	$k=3$	$k=3$	$k=2$	\mathcal{M}_2	1.28	$k=2$	$k=1$	$k=3$	$k=2$
\mathcal{M}_3	5.04	$k=2$	$k=3$	$k=2$	$k=2$	\mathcal{M}_3	1.55	$k=2$	$k=3$	$k=1$	$k=2$
\mathcal{M}_4	325.17	$k=2$	$k=1$	$k=2$	$k=2$	\mathcal{M}_4	1.79	$k=2$	$k=3$	$k=3$	$k=2$
\mathcal{M}_5	1235.60	$k=2$	$k=1$	$k=1$	$k=2$	\mathcal{M}_5	4.62	$k=2$	$k=1$	$k=2$	$k=2$

Note. $\mathcal{M}_1^{\text{HPM}}$ is the highest probability model. The posterior odds, for $\mathcal{M}_i, i = 2, 3, 4$, was computed as the ratio of model probabilities $Pr(\mathcal{M}_1) / Pr(\mathcal{M}_i)$. This corresponds to the Bayes factor BF_{1i} because each model had equal prior probabilities. k is the respective mixture component (Figure 1, panel A)