

A note on the normality assumption for Bayesian models of constraint in behavioral individual differences

Thomas J. Faulkenberry

Tarleton State University, Department of Psychological Sciences, Stephenville, TX, United States

Abstract

To investigate the structure of individual differences in performance on behavioral tasks, Haaf and Rouder (2017) developed a class of hierarchical Bayesian mixed models with varying levels of constraint on the individual effects. The models are then compared via Bayes factors, telling us which model best predicts the observed data. One common criticism of their method is that the observed data are assumed to be drawn from a normal distribution. However, for most cognitive tasks, the primary measure of performance is a response time, the distribution of which is well known not to be normal. In this paper, I investigate the assumption of normality for two datasets in numerical cognition. Specifically, I show that using a shifted lognormal model for the response times does not change the overall pattern of inference. Further, since the model-estimated effects are now on a logarithmic scale, the interpretation of the modeling becomes more difficult, particularly because the estimated effect is now multiplicative rather than additive. As a result, I recommend that even though response times are not normally distributed in general, the simplification afforded by the Haaf and Rouder approach provides a pragmatic approach to modeling individual differences in behavioral tasks.

Keywords: Bayesian mixed models, hierarchical models, Bayes factors, response time models, log-normal distribution

*Corresponding author

Email address: faulkenberry@tarleton.edu (Thomas J. Faulkenberry)

ORCID iD: [0000-0002-8997-3794](https://orcid.org/0000-0002-8997-3794) (Thomas J. Faulkenberry)

1. Introduction

In the behavioral sciences, a common target of investigation is individual performance on behavioral tasks, and particularly whether this individual performance can predict other measureable outcomes. As an illustrative example, consider a simple number comparison task where subjects are asked to choose the physically larger of two number digits presented in different physical sizes on a screen (e.g., a large numeral 2 displayed alongside a small numeral 8). Performance on this task routinely exhibits a *size congruity effect* (Henik & Tzelgov, 1982), where people are slower (on average) to choose the larger when the numbers are presented in a physical size configuration that is incongruent with their relative numerical magnitude. Importantly, researchers often use individual performance on this task to predict other meaningful behavioral outcomes, especially those related to mathematics anxiety and ability. For example, Rubinsten and Henik (2005) found that people with developmental dyscalculia exhibited a smaller size congruity effect compared to a typical-functioning control group, which they interpreted as evidence for a lack of automatic number activation in the dyscalculia group.

Given that individual performance on behavioral tasks is used in this metric sense to predict other tangible outcomes, a natural question concerns whether this metric scale has any constraint. Specifically, Haaf and Rouder (2017) proposed that a method to ascertain whether people differ in performance solely in a *quantitative* fashion (i.e., everybody exhibits the effect in the same direction, but differ in the size of the effect), or whether there are also *qualitative* individual differences, where some people exhibit a positive effect, but others exhibit the effect in the opposite direction. Such questions are becoming important in the psychological and behavioral sciences, particularly in terms of providing much-needed constraint on the plethora of observed “effects” and helping to guide more targeted theoretical development about human behavior (Rouder & Haaf, 2021).

The purpose of this paper is to examine one of the fundamental statistical assumptions of the Haaf and Rouder (2017) method for modeling constraint on behavioral individual differences. In brief, the method relies on assuming that the resultant behavioral measures (e.g., response times) are drawn from a normal distribution whose mean is represented as a linear combination of a variable intercept and slope (effect). In turn, each of these parameters is further drawn from normal distributions centered at 0 and scaled according to overall variability. Different models of individual difference structure are instantiated by placing varying levels of constraint on the slope/effect parameter. Critical to the Haaf and Rouder (2017) method is a Bayesian comparison of these models, which uses a combination of the well-known analysis of variance approach developed by Rouder et al. (2012) and the encompassing prior approach (Faulkenberry, 2019). The approach has been used successfully to investigate the structure of individual differences in many behavioral phenomena, including Stroop and Simon effects (Haaf & Rouder, 2019), the truth effect (Schnuerch et al., 2020), numerical distance effects (Vogel et al., 2021), and the numerical size congruity effect (Faulkenberry & Bowman, 2020).

One criticism of the Haaf and Rouder (2017) method is the assumption that raw performance measures are drawn from a normal distribution. This criticism is particularly salient when the primary measure is response time, as response times are well known to exhibit a distinct positive skew. While there are many methods for modeling response times using skewed distributions (i.e., ex-Gaussian, inverse Gaussian/Wald, etc.), the implementation of such distributions into the Haaf and Rouder (2017) framework is quite difficult. Compared to a normal distribution, these distributions involve multiple parameters, and so it is not clear on which parameter the congruity effect should be applied. Further, the Haaf and

Rouder (2017) method is specifically built on an established computational framework built into the BayesFactor package in R, so deviating from this method would require the user to develop many new techniques from first principles. As an alternative, one simple approach that might prove attractive is to assume that the observed response times follow a (shifted) lognormal distribution; then, the analyst may simply transform the observed response times by first shifting by a fixed amount (e.g., 200 ms is a common recommendation) and then taking the (natural) logarithm. The resulting distribution of (log) response times is then approximately normal and may be “fed into” the Haaf and Rouder (2017) method with little difficulty.

The purpose of this brief paper is to investigate this approach and to argue two points:

1. the inferences obtained from the shifted lognormal model are practically equivalent as those obtained from the original model with the normality assumption; and
2. interpreting the estimated model parameters from the shifted lognormal model is nontrivial and potentially inappropriate in the context of these behavioral tasks.

2. Bayesian model implementation

First, I will describe the Bayesian mixed model approach developed by Haaf and Rouder (2017), particularly as applied to behavioral tasks where the primary observed data are response times. Before going into the details, I will reiterate that the main aim of this approach is to build a (hypothetical) generative process for each observed response time in a behavioral task. That is, there is no aggregation of trials at the individual or group level that needs to occur.

Each observed response time is assumed to be the sum of four components: (i) a grand mean μ ; (ii) a subject-specific adjustment α to the grand mean (i.e., so that $\mu + \alpha$ gives a “random” intercept for each subject); (iii) a subject-specific effect term δ ; and (iv) a noise term ε . The hierarchical model is then built by assuming each of these components is drawn from some to-be-defined probability distribution. Of particular interest is the distribution that generates each subject’s effect term—this distribution is the one on which we build our competing models of individual difference structure.

We let Y_{ijk} denote the response time for the k th replicate of the i th subject in the j th experimental condition (usually two conditions, so that $j = 1, 2$). As described, our random effects linear model on the vector of response times Y_{ijk} looks like:

$$Y_{ijk} \sim \mathcal{N}(\mu + \alpha_i + x_j \cdot \delta_i, \sigma^2).$$

Here, μ denotes the grand mean intercept and α_i represents the specific intercept adjustment for subject i . The term x_j is a binary variable which codes the experimental condition for each trial. For example, suppose we are interested in modeling a congruity effect, where response times on incongruent trials generally increase compared to those of congruent trials. In this case, for congruent trials (condition $j = 1$), we would set $x_1 = 0$, and for incongruent trials (condition $j = 2$), we would set $x_2 = 1$. Under such a specification, δ_i then represents the (random) congruity effect for subject i . Finally, σ^2 represents the latent sampling variance of the observed response times.

The next step is to propose a structure for the parent distribution of random effects δ_i (i.e., the distribution from which each subject’s size-congruity effect δ_i is randomly drawn). We define four possible populations for these δ_i , each of which mathematically specifies one of four possible theoretical positions about the distribution of effects.

2.1. The unconstrained model

The unconstrained model, denoted \mathcal{M}_u , allows the effects δ_i to vary both in type/quality (i.e., positive or negative) as well as magnitude. As such, with \mathcal{M}_u we place no constraint on the individual effects δ_i . We define this model as

$$\mathcal{M}_u : \delta_i \sim \mathcal{N}(v, \eta^2),$$

where v and η^2 represent the mean and variance, respectively, of the distribution of individual effects δ_i .

2.2. The positive-effects model

The positive-effects model, denoted \mathcal{M}_+ , hypothesizes that effects δ_i only vary in quantity (i.e., they are always positive, but possibly differ in magnitude between subjects). \mathcal{M}_+ is a constrained model in the sense that it specifies the assumption that all individual effects δ_i are positive. That is,

$$\mathcal{M}_+ : \delta_i \sim \mathcal{N}_+(v, \eta^2),$$

where \mathcal{N}_+ denotes a truncated normal distribution with lower bound 0.

2.3. The common-effect and null models

Whereas the unconstrained and positive-effects models are usually the primary players in studies on individual structure, the common-effect and null models are defined to provide a critical check of experimental design. The common-effect model places even more constraint on the distribution of effects by assuming that each individual has the *same* effect. That is,

$$\mathcal{M}_1 : \delta_i = v,$$

Such a model serves to probe the following question: if the common-effect model was the best predictor of the observed data, one would be forced to question the efficiency of the experimental design as a test to elicit individual differences in the effect. As one might expect, the null model is the most constrained of the four, as it specifies that each subject's size-congruity effect is zero:

$$\mathcal{M}_0 : \delta_i = 0.$$

It is used for a similar reason: if the null model was the best predictor of the observed data, then one must question the efficiency of the experimental design to elicit effects of any sort.

2.4. Prior specifications

Generally, most applications of the Haaf and Rouder (2017) method follow similar "default" prior specifications. The critical parameters I'll describe here are δ_i , v , and η^2 . The default procedure is to use the g-prior approach (Rouder et al., 2012; Zellner, 1986), which re-expresses these parameters as a standardized effect size. To see how this works, consider the collection of individual effect parameters δ_i . We define $g_\delta = \eta^2/\sigma^2$, yielding a hyperparameter that casts the variability of δ_i in terms of the ratio of true variability η^2 to sampling variability σ^2 . With this we can re-write our unconstrained model as

$$\mathcal{M}_u : \delta_i \sim \mathcal{N}(v, g_\delta \sigma^2).$$

Similarly, we may scale the mean size-congruity effect v in terms of sampling variability and get a new hyperparameter g_v . Continuing up the hierarchy, these new (hyper)parameters

need priors as well. The default specification (Zellner, 1986) is to use scaled inverse- χ^2 distributions with one degree of freedom and scale r^2 .

To be clear, the g -prior setup is quite clever, as it completely describes these critical parameters in terms of sampling variability σ^2 . By doing this, we convert the problem of specifying priors on δ_i , ν , and η^2 into one where we simply need to specify the expected variability of our effects relative to the expected overall variability of the observed response times. Like Haaf and Rouder (2017), I will use $\sigma = 300$ ms as a prior expectation for the variability of observed response times.

Now we can actually finish setting our priors. First we consider g_ν , the g -prior on the mean size-congruity effect. With the g -prior setup, we assume that $\nu \sim \mathcal{N}(0, g_\nu \sigma^2)$, where $g_\nu \sim \text{Scale-inv-}\chi^2(r_\nu^2)$. The scale parameter r_ν should reflect our prior belief about the relative magnitude of our expected effects. For the types of effects we often see in numerical cognition (and certainly the types of tasks we will describe in this paper), I usually expect such effects to be, on average, around 50 ms, or 1/6 of the expected overall trial-by-trial variability ($\sigma = 300$ ms). Thus, we set $r_\nu = 1/6$.

Second, we consider g_δ , which describes the variability of individual effects around the mean effect. With the g -prior setup, we assume that $g_\delta \sim \text{Scale-inv-}\chi^2(r_\delta^2)$. Like Haaf and Rouder (2017), we set $r_\delta = 1/10$, which would indicate that the expected variability of the effect across individuals should be about 1/10 of $\sigma = 300$ ms, or around 30 ms.

2.5. Model comparison

Since our goal is to capture the latent structure of individual differences in the effects we observe in our behavioral task, our problem is first and foremost one of *model comparison*. That is, we ask which of the four competing models defined above is the most adequate as a predictor of our observed data? To answer this question, we use Bayes factors (Jeffreys, 1961; Kass & Raftery, 1995), which index the relative predictive adequacy of two models by comparing the marginal likelihood of observed data under one model to another (Faulkenberry, 2022; Faulkenberry et al., 2020). For example, a Bayes factor of 10 indicates that the observed data are 10 times more likely under one model compared to another. Techniques for computing Bayes factors among three of the four models above (\mathcal{M}_u , \mathcal{M}_1 , \mathcal{M}_0) were previously developed by Rouder et al. (2012) and are implemented in the BayesFactor (Morey et al., 2018) package in R (R Core Team, 2022). The Bayes factor between the constrained positive effects model \mathcal{M}_+ and the unconstrained model \mathcal{M}_u is computed by the *encompassing prior* method (Faulkenberry, 2019; Klugkist et al., 2005), which is based on counting the number of posterior samples of \mathcal{M}_u which obey the constraint placed by \mathcal{M}_+ , then comparing this to the number of prior samples which obey the same constraint.

3. Case studies

My goal in this paper is to compare the inferences from the default Haaf and Rouder (2017) method, which assumes that the observed response times are drawn from a *normal distribution*, to a modified approach where the observed response times are drawn from a *lognormal distribution*. To do this, I will perform two case studies where I analyze two datasets that have already appeared in the literature. In Case study 1, I will model the latent structure of individual differences in the *size congruity effect* (Henik & Tzelgov, 1982), a classic phenomenon in numerical cognition in which people are slower to choose the larger of two presented numbers when the numbers are presented in a physical size that is incongruent with their relative numerical magnitude (e.g., a large numeral 2 displayed

alongside a small numeral 8). The data for Case study 1 (19 499 response times from $N = 53$ subjects) were originally reported in Faulkenberry and Bowman (2020). In Case study 2, I will model the latent structure of individual differences in the *unit decade compatibility effect*, another classic phenomenon in numerical cognition (Nuerk et al., 2001). The data for Case study 2 (11 600 response times from $N = 53$ subjects) are unpublished but available as part of a collaborative pre-registration project by Cipora et al. (2021).

3.1. Case study 1—size congruity effect

The first analysis I will describe is the default Haaf and Rouder (2017) method, which places a normal distribution on the observed response times. The individual effect estimates from the unconstrained model are displayed in the left column of Figure 1. We can see that the observed effects for each subject (denoted by black crosses) span from -14.59 ms to 142.10 ms. In this context, we compute observed effects by subtracting each subject’s mean response time for congruent trials from the mean response time for incongruent trials. With the exception of one subject, the observed size-congruity effects are all constrained to be positive. Estimates from the hierarchical Bayesian model are displayed as blue dots with shaded 95 % credible intervals. These estimates are computed as means of the posterior samples for each δ_i , and the 95 % credible intervals are computed as the central 95 % of the posterior samples (i.e., ranging between the 2.5 % and 97.5 % quantiles of the samples). The red dashed line represents an (posterior) estimated mean effect of $\nu = 60$ ms.

As is usually seen with this type of modeling (and hierarchical modeling in general), we observe a fair amount of *shrinkage* in our estimates. Notice that the estimated effects (the blue dots) extend over a smaller range (8.84 ms to 115.73 ms) than the observed effects (the black crosses; -14.59 ms to 142.10 ms). This shrinkage reflects how the hierarchical model accounts for sampling variability at all levels.

The right column of Figure 1 shows the Bayes factor model comparisons. As we can see, the observed data were 7.19 times more likely under the positive-effects model \mathcal{M}_+ than under the unconstrained model \mathcal{M}_u . If we assume 1-to-1 prior odds for \mathcal{M}_+ and \mathcal{M}_u , this means that our posterior odds in favor of \mathcal{M}_+ have increased to 7.19-to-1, which is equivalent to a posterior probability of $\Pr(\mathcal{M}_+ \mid \text{data}) = 0.88$. These models were overwhelmingly preferred over the common-effect model \mathcal{M}_1 and the null model \mathcal{M}_0 , as \mathcal{M}_u was more likely to have predicted the observed data by factors of 10^{11} -to-1 and 10^{156} -to-1, respectively.

Next, we perform the same procedure while assuming a shifted lognormal distribution on the observed response times. To do this, we transform the observed response times by first subtracting a constant amount from each response time (here, I chose a shift of 200 ms), then taking the (natural) logarithm of the result. As we can see in Figure 2, the transformed distribution appears approximately normal, indicating that the lognormal model is appropriate in this case study as well. The resulting transformed data can be directly modeled as above, the results of which I will now describe.

The overall similarity of these results with the first analysis is striking. We see very similar patterns of observed effects, estimated effects, and shrinkage. For the log transformed data, we see a posterior estimated common effect (red dashed line) $\nu = 0.15$. If we back-transform this back to the original response time scale, we get an estimated common effect of 1.16. Because the data are on a logarithmic scale, this effect is multiplicative, so an estimated effect of 1.16 is a 16 % increase in response times. For these data, this is roughly equivalent to a response time increase of 86 ms.

The similarity persists with the Bayes factor comparisons. In the right column of




Figure 1: Individual effect estimates (left column) and Bayes factor model comparisons (right column) for Case study 1 under a normal distribution assumption. Posterior means and 95 % credible intervals for δ_i are represented by blue dots and a gray band, respectively. The + symbols represent the observed size-congruity effect for each subject. The red dashed-line represents the estimated mean size-congruity effect ν . For the model comparisons, the red box denotes the winning model, and Bayes factors are displayed beside each arrow.




Figure 2: Distributions of observed response times in the size congruity task (Case study 1). The left panel displays the original observed response times, whereas the right panel displays the log-transformed response times.

Figure 3 we can see the observed data were 6.21 times more likely under the positive-effects model \mathcal{M}_+ than under the unconstrained model \mathcal{M}_u . Further, these models were again overwhelmingly preferred over the common-effect model \mathcal{M}_1 and the null model \mathcal{M}_0 .




Figure 3: Individual effect estimates (left column) and Bayes factor model comparisons (right column) for Case study 1 under a *lognormal* distribution assumption. Posterior means and 95 % credible intervals for δ_i are represented by blue dots and a gray band, respectively. The + symbols represent the observed size-congruity effect for each subject. The red dashed-line represents the estimated mean size-congruity effect ν . For the model comparisons, the red box denotes the winning model, and Bayes factors are displayed beside each arrow.

In all, it seems that with the exception of the raw effect estimate, the inferences we obtain from using a shifted lognormal model on observed response times is very similar to that when we use the default normal specifications recommended by Haaf and Rouder (2017). In both cases, the positive effects model is preferred over the unconstrained model.

3.2. Case study 2—unit decade compatibility effect

As above, I will first report the results of modeling using the default Haaf and Rouder (2017) method with a normal distribution on the observed response times. The individual effect estimates from the unconstrained model are displayed in the left column of Figure 4. The observed effects for each subject (denoted by black crosses) span from -5.37 ms to 137.52 ms. Similar to Case study 1, the observed effects were mostly positive. Estimates from the hierarchical Bayesian model are displayed as blue dots with shaded 95 % credible interval. The red dashed line represents an (posterior) estimated mean effect of $\nu = 43$ ms. Note that we again observe shrinkage in our estimates, as the estimated effects extend from 15.04 ms to 92.80 ms), a smaller range than that of the observed estimates.

The right column of Figure 4 shows the Bayes factor model comparisons. In this case, the observed data were 4.17 times more likely under the positive-effects model \mathcal{M}_+ than under the unconstrained model \mathcal{M}_u . If we assume 1-to-1 prior odds for \mathcal{M}_+ and \mathcal{M}_u , this means that our posterior odds in favor of \mathcal{M}_+ have increased to 4.17-to-1, which is equivalent to a posterior probability of $\Pr(\mathcal{M}_+ | \text{data}) = 0.81$. As in Case study 1, these models were strongly preferred over the common-effect model \mathcal{M}_1 and the null model \mathcal{M}_0 .

Next, we run the analysis again, but this time assuming a shifted lognormal distribution on the observed response times. As before, we transform the observed response times




Figure 4: Individual effect estimates (left column) and Bayes factor model comparisons (right column) for Case study 2 under a normal distribution assumption. Posterior means and 95 % credible intervals for δ_i are represented by blue dots and a gray band, respectively. The + symbols represent the observed size-congruity effect for each subject. The red dashed-line represents the estimated mean size-congruity effect ν . For the model comparisons, the red box denotes the winning model, and Bayes factors are displayed beside each arrow.

by first subtracting a constant amount from each response time (here, I chose a shift of 200 ms), then taking the (natural) logarithm of the result. As we can see in Figure 5, the transformed distribution appears approximately normal, indicating that the lognormal model is appropriate here.




Figure 5: Distributions of observed response times in the numerical comparison task (Case study 1). The left panel displays the original observed response times, whereas the right panel displays the log-transformed response times.

As with Case study 1, we see very similar patterns of observed effects, estimated effects, and shrinkage in Figure 4. For the log transformed data, we see a posterior estimated common effect (red dashed line) $\nu = 0.08$. On the original response time scale, this is equivalent to an estimated (multiplicative) common effect of 1.09, a 9 % increase in response times. For these data, this is roughly equivalent to a response time increase of 57 ms.

The Bayes factor comparisons also present the same message. In the right column of Figure 6 we can see the observed data were 8.40 times more likely under the positive-effects model \mathcal{M}_+ than under the unconstrained model \mathcal{M}_u . Further, these models were again preferred over the common-effect model \mathcal{M}_1 and the null model \mathcal{M}_0 . Again, the inference from using a shifted lognormal model on observed response times is very similar to that when we use the default normal specifications recommended by Haaf and Rouder (2017). In both cases, the positive effects model is preferred over the unconstrained model.

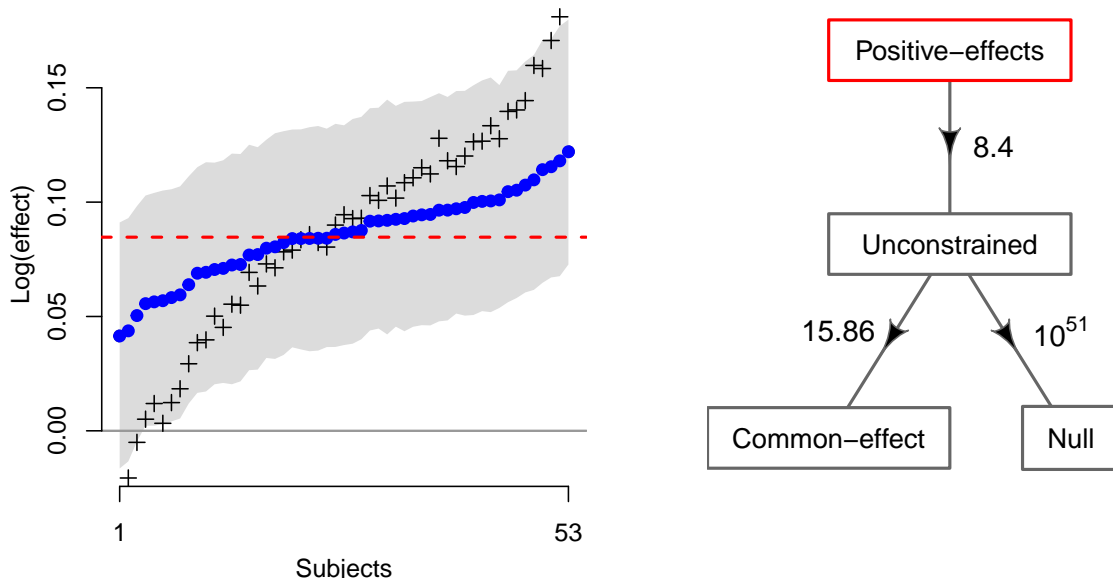


Figure 6: Individual effect estimates (left column) and Bayes factor model comparisons (right column) for Case study 2 under a *lognormal* distribution assumption. Posterior means and 95% credible intervals for δ_i are represented by blue dots and a gray band, respectively. The + symbols represent the observed size-congruity effect for each subject. The red dashed-line represents the estimated mean size-congruity effect ν . For the model comparisons, the red box denotes the winning model, and Bayes factors are displayed beside each arrow.

4. Simulation study

As we can see from the previous two sections, both case studies lead to a common conclusion. Even though the observed response times exhibit positive skew, the inference we obtain from applying the default Haaf and Rouder (2017) method (which assumes a normal distribution on response times) is practically the same as when we apply a shifted lognormal model on response times. To extend support for this tentative conclusion, I performed a simulation study to benchmark and compare the long-term performance of both methods against data which are assumed to be generated from either the positive-effects model or the unconstrained model. In this section, I describe the simulation study and report its results. The simulation was performed in R, and the simulation script can be viewed at <https://bit.ly/3D9QGIZ>.

For each simulation run, the data were assumed to be generated from a hierarchical shifted-Wald distribution (e.g., Anders et al., 2016; Faulkenberry et al., 2018). In general, a shifted-Wald distribution represents the collection of stopping times for a continuous accumulator with drift toward a fixed response boundary. The distribution exhibits a positive skew characteristic of response time distributions, which positions it as a good model for

use here. The shifted-Wald can be completely described by three parameters: drift rate γ , which represents the rate at which information is accumulated during stimulus presentation; response threshold α , which represents the amount of information that must be accumulated before a decision can be initiated; and shift θ , which represents the remaining portion of the response time on a trial which is not accounted for by the accumulation process (i.e., perceptual encoding, motor preparation, etc.).

To explain how the data were simulated, consider the context of a typical repeated-measures experiment with an equal number of congruent and incongruent trials. Let us assume that N subjects each produce k trials in the congruent condition. The k “observed” RTs for each subject $i = 1, \dots, N$ are assumed to be randomly drawn from a subject-specific shifted-Wald distribution with drift rate γ_i , response threshold α_i , and shift θ_i . The collection of parameters γ_i , α_i , and θ_i for $i = 1, \dots, N$ were each randomly drawn from normal distributions with mean and variance set to match the shifted-Wald parameter estimates given in Faulkenberry et al. (2018). That is:

$$\begin{aligned}\gamma_i &\sim \mathcal{N}(3.91, 0.70^2) \\ \alpha_i &\sim \mathcal{N}(0.92, 0.17^2) \\ \theta_i &\sim \mathcal{N}(0.32, 0.05^2).\end{aligned}$$

To produce the k incongruent trials for each subject i , I followed the same sampling scheme, but additionally instantiated a congruity effect δ_i on the shift parameter θ_i for each subject $i = 1, \dots, N$. Thus the observed incongruent trial RTs for subject i were generated from a shifted-Wald distribution with drift rate γ_i , response threshold α_i , and shift $\theta_i + \delta_i$. Individual differences in the congruity effects were generated following the method of Rouder and Haaf (2019), who assumed that the subject-level effects δ_i were drawn in a hierarchical fashion as $\delta_i \sim \mathcal{N}(\mu_\delta, \sigma_\delta^2)$, where we must further specify our assumptions on μ_δ and σ_δ^2 . It is with these assumptions that we can specify the two models which constrain the individual effects δ_i . For data generated under the unconstrained model \mathcal{M}_u , the overall mean effect μ_δ was drawn as $\mu_\delta \sim \mathcal{N}(0.1, 0.1^2)$. For data generated under the positive-effects model \mathcal{M}_+ , μ_δ was assumed to be drawn from a truncated normal: $\mu_\delta \sim \mathcal{N}_+(0.1, 0.05^2)$. For both models, we assumed $\sigma_\delta^2 \sim \text{Inv-Gamma}(2, 0.03^2)$.

On each simulation run, the collection of $2 \cdot k \cdot N$ RTs were submitted to each of the modeling workflows described earlier in the paper. In the first workflow, I applied the default method of Haaf and Rouder (2017), which assumes the observed RTs are generated via a normal distribution. In the second workflow, I transformed the observed RTs by first subtracting 0.95 times the minimum observed RT (i.e., shifting the distribution to remove the leading edge) and then taking the natural logarithm of the shifted RTs. In all, I completed 200 simulation runs in each of 6 conditions created by systematically varying N and k to reflect common experimental designs in the cognitive and behavioral sciences. Specifically, I crossed $N = 20, 80$ with $k = 50, 100, 200$.

First, let us consider the accuracy of the default and log-transform methods. Tables 1 and 2 depict the accuracies obtained from the two methods for each of our 6 experimental conditions for data generated under the positive-effects model \mathcal{M}_+ and the unconstrained model \mathcal{M}_u , respectively. For data generated under \mathcal{M}_+ , inference was very accurate for conditions with $N = 20$ subjects. Even with smaller numbers of trials (e.g., $k = 50$), both the default and log-transform methods were largely correct in their model choices, and there was no obvious difference between the two methods. For larger numbers of subjects ($N = 80$), the results were curious. The log-transform method was consistently more

accurate than the default method of Haaf and Rouder (2017), with an accuracy advantage between 5 % and 9 %. Both methods exhibited smallest accuracy when $k = 50$, but these accuracies increased with increasing k . For data generated under \mathcal{M}_u , performance was consistently at ceiling. Curiously, the worst performance (accuracy = 97 %) occurred in the condition with a small number of participants ($N = 20$) each contributing a large number of trials ($k = 200$).

Table 1: Model choice accuracy for the default Haaf and Rouder (2017) method and the log-transform method, calculated as the proportion of datasets simulated under the positive-effects model \mathcal{M}_+ for which $\text{BF}_{+u} > 1$

k	$N = 20$		$N = 80$	
	Default	Log-transform	Default	Log-transform
50	0.93	0.93	0.51	0.56
100	0.92	0.93	0.52	0.61
200	0.95	0.97	0.71	0.79

Table 2: Model choice accuracy for the default Haaf and Rouder (2017) method and the log-transform method, calculated as the proportion of datasets simulated under the positive-effects model \mathcal{M}_+ for which $\text{BF}_{+u} > 1$

k	$N = 20$		$N = 80$	
	Default	Log-transform	Default	Log-transform
50	0.99	0.99	0.99	0.99
100	0.99	0.99	0.99	0.99
200	0.97	0.97	0.99	0.99

While accuracy of the methods is important to assess, the critical claim of this paper is that the default workflow of Haaf and Rouder (2017) produces the same inference as the log-transform method. We can assess this claim empirically by considering the consistency of the inferences obtained by both methods in our simulated datasets. Table 3 shows that these methods exhibit a great deal of consistency. For data generated under \mathcal{M}_+ , model choice consistency was extremely high (at least 95 %) for conditions with $N = 20$. Similar to the results with accuracy above, model choice consistency was not quite as high for conditions with $N = 80$, though consistency did increase with increasing k . For data generated under \mathcal{M}_u , model choice consistency was at ceiling for all conditions.

5. Conclusion

The main aim of this paper was to compare the inferences from two methods for assessing the structure of individual differences in behavioral tasks. The first of these two methods was the default Haaf and Rouder (2017) method, which assumes that the observed response times are drawn from a normal distribution. The second was a modified approach where the observed response times are assumed to follow a lognormal distribution. Two case studies

Table 3: Model choice consistency for the default Haaf and Rouder (2017) method and the log-transform method, calculated as the proportion of datasets for which both methods choose the same model

k	Positive-effects		Unconstrained	
	$N = 20$	$N = 80$	$N = 20$	$N = 80$
50	0.95	0.83	0.99	0.99
100	0.95	0.85	0.99	0.99
200	0.98	0.90	0.99	0.99

and a simulation lead to a common conclusion. Even though observed response times typically exhibit positive skew, the inference we obtain from applying the default Haaf and Rouder (2017) method is practically equivalent to those obtained when we apply a shifted lognormal model to the response times. As we saw in both case studies, applying a shift and then taking the natural logarithm of the observed response times does indeed transform the distribution of observed data into one which is approximately normal. Certainly, the Haaf and Rouder (2017) method works well for this transformed data, but the penalty is in the interpretation. When the observed data is transformed to the log scale, the “effects” we see in the data (i.e., differences between the observed data that occur as a function of the experimental manipulation) are now differences in the log scale. Differences in the log scale become multiplicative differences (i.e., quotients) when we transform back to the original scale of the response times. While multiplicative effects can make sense in many contexts, such effects are not typical in the context of effects on response time. Indeed, most typical response time models assume that total response time is the *sum* of its constituent subprocesses (Ashby & Townsend, 1980; Schwarz, 2001). As such, it is not clear how one of these behavioral or cognitive effects could reasonably be interpreted in a multiplicative context.

Given that (i) the pattern of inference does not change, and (ii) the interpretation of estimated effects becomes less clear, there is no compelling reason to reject the normal assumption on response times when applying the Haaf and Rouder (2017) method for investigating individual difference structures in behavioral tasks.

References

- Anders, R., Alario, F.-X., & Maanen, L. V. (2016). The shifted Wald distribution for response time data analysis. *Psychological Methods*, 21(3), 309–327. <https://doi.org/10.1037/met0000066>
- Ashby, F., & Townsend, J. T. (1980). Decomposing the reaction time distribution: Pure insertion and selective influence revisited. *Journal of Mathematical Psychology*, 21(2), 93–123. [https://doi.org/10.1016/0022-2496\(80\)90001-2](https://doi.org/10.1016/0022-2496(80)90001-2)
- Cipora, K., Faulkenberry, T. J., Bahnmüller, J., Connolly, H., Bowman, K., Moeller, K., & Nuerk, H.-C. (2021). *Prevalence of cognitive phenomena: Comparison of four methods* [Project Homepage]. Center for Open Science. <https://osf.io/9y6m4/>
- Faulkenberry, T. J. (2019). A tutorial on generalizing the default Bayesian t-test via posterior sampling and encompassing priors. *Communications for Statistical Applications and Methods*, 26(2), 217–238. <https://doi.org/10.29220/csam.2019.26.2.217>

- Faulkenberry, T. J. (2022). *Psychological statistics*. Routledge. <https://doi.org/10.4324/9781003181828>
- Faulkenberry, T. J., & Bowman, K. (2020). *Modeling the latent structure of individual differences in the numerical size-congruity effect*. PsyArXiv. <https://doi.org/10.31234/osf.io/4b9rs>
- Faulkenberry, T. J., Ly, A., & Wagenmakers, E.-J. (2020). Bayesian inference in numerical cognition: A tutorial using JASP. *Journal of Numerical Cognition*, 6(2), 231–259. <https://doi.org/10.5964/jnc.v6i2.288>
- Faulkenberry, T. J., Vick, A. D., & Bowman, K. A. (2018). A shifted Wald decomposition of the numerical size-congruity effect: Support for a late interaction account. *Polish Psychological Bulletin*, 49(4), 391–397. <https://doi.org/10.24425/119507>
- Haaf, J. M., & Rouder, J. N. (2017). Developing constraint in Bayesian mixed models. *Psychological Methods*, 22(4), 779–798. <https://doi.org/10.1037/met0000156>
- Haaf, J. M., & Rouder, J. N. (2019). Some do and some don't? Accounting for variability of individual difference structures. *Psychonomic Bulletin & Review*, 26, 772–789. <https://doi.org/10.3758/s13423-018-1522-x>
- Henik, A., & Tzelgov, J. (1982). Is three greater than five: The relation between physical and semantic size in comparison tasks. *Memory & Cognition*, 10(4), 389–395. <https://doi.org/10.3758/bf03202431>
- Jeffreys, H. (1961). *The theory of probability* (3rd ed.). Oxford University Press.
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90(430), 773–795. <https://doi.org/10.2307/2291091>
- Klugkist, I., Kato, B., & Hoijsink, H. (2005). Bayesian model selection using encompassing priors. *Statistica Neerlandica*, 59(1), 57–69. <https://doi.org/10.1111/j.1467-9574.2005.00279.x>
- Morey, R. D., Rouder, J. N., Jamil, T., Urbanek, S., Forner, K., & Ly, A. (2018). *BayesFactor: Computation of Bayes factors for common designs* (Version 0.9.12-4.2) [Computer software]. The Comprehensive R Archive Network. <https://CRAN.R-project.org/package=BayesFactor>
- Nuerk, H.-C., Weger, U., & Willmes, K. (2001). Decade breaks in the mental number line? Putting the tens and units back in different bins. *Cognition*, 82(1), B25–B33. [https://doi.org/10.1016/s0010-0277\(01\)00142-1](https://doi.org/10.1016/s0010-0277(01)00142-1)
- R Core Team. (2022). *R: A language and environment for statistical computing* (Version 4.2.1) [Computer software]. R Foundation for Statistical Computing. <https://www.r-project.org/>
- Rouder, J. N., & Haaf, J. M. (2019). *Why most studies of individual differences with inhibition tasks are bound to fail*. PsyArXiv. <https://doi.org/10.31234/osf.io/3cjr5>
- Rouder, J. N., & Haaf, J. M. (2021). Are there reliable qualitative individual difference in cognition? *Journal of Cognition*, 4(1), 46. <https://doi.org/10.5334/joc.131>
- Rouder, J. N., Morey, R. D., Speckman, P. L., & Province, J. M. (2012). Default Bayes factors for ANOVA designs. *Journal of Mathematical Psychology*, 56(5), 356–374. <https://doi.org/10.1016/j.jmp.2012.08.001>
- Rubinsten, O., & Henik, A. (2005). Automatic activation of internal magnitudes: A study of developmental dyscalculia. *Neuropsychology*, 19(5), 641–648. <https://doi.org/10.1037/0894-4105.19.5.641>
- Schnuerch, M., Nadarevic, L., & Rouder, J. N. (2020). The truth revisited: Bayesian analysis of individual differences in the truth effect. *Psychonomic Bulletin & Review*, 28(3), 750–765. <https://doi.org/10.3758/s13423-020-01814-8>

- Schwarz, W. (2001). The ex-Wald distribution as a descriptive model of response times. *Behavior Research Methods, Instruments, & Computers*, 33(4), 457–469. <https://doi.org/10.3758/bf03195403>
- Vogel, S. E., Faulkenberry, T. J., & Grabner, R. H. (2021). Quantitative and qualitative differences in the canonical and the reverse distance effect and their selective association with arithmetic and mathematical competencies. *Frontiers in Education*, 6. <https://doi.org/10.3389/feduc.2021.655747>
- Zellner, A. (1986). On assessing prior distributions and Bayesian regression analysis with g-prior distributions. In P. K. Goel & A. Zellner (Eds.), *Bayesian inference and decision techniques: Essays in honor of Bruno de Finetti* (pp. 233–243). North-Holland. <https://doi.org/10.1037/14731-006>