Consequences of Using Power Transforms as a Statistical Solution in Linear Mixed-Effects

Models of Chronometric Data

Van Rynald T. Liceralde & Peter C. Gordon

Department of Psychology and Neuroscience, University of North Carolina at Chapel Hill

Author Note

We thank Melvin Yap for providing the raw individual-difference and trial-level files from the English Lexicon Project; Cindy Vo for assisting in preprocessing said files; Patrick Curran, Jennifer Arnold, Daniel Bauer, Giulia Pancani, Fanting Kung, and the Language, Cognition, and Brain lab for helpful discussions on this work. This project was based on VRTL's Master's thesis. The preprint, data, and analysis scripts for this project may be found at <https://osf.io/vyu65/>.

Address correspondence to:

Van Rynald T. Liceralde
Department of Psychology and Neuroscience, CB 3270
University of North Carolina at Chapel Hill
Chapel Hill, NC 27599-3270
vliceral@live.unc.edu

12,494 words and numbers

Abstract

Power transforms have been increasingly used in linear mixed-effects models (LMMs) of chronometric data (e.g., response times [RTs]) as a statistical solution to preempt violating the assumption of residual normality. However, differences in results between LMMs fit to raw RTs and transformed RTs have reignited discussions on issues concerning the transformation of RTs. Here, we analyzed two word-recognition megastudies and performed Monte Carlo simulations to better understand the consequences of transforming RTs in LMMs. Within each megastudy, transforming RTs produced different fixed- and random-effect patterns; across the megastudies, RTs were optimally normalized by different power transforms, and results were more consistent among LMMs fit to raw RTs. Moreover, the simulations showed that LMMs fit to optimally normalized RTs had greater power for main effects in smaller samples, but that LMMs fit to raw RTs had greater power for interaction effects as sample sizes increased, with negligible differences in Type I error rates between the two models. Based on these results, we recommend that LMMs be fit to raw RTs when there is no compelling reason beyond nonnormality to transform RTs and when the interpretive framework mapping the predictors and RTs treats RT as an interval scale.

Keywords: linear mixed-effects models, log and inverse transform, normality, response times, Monte Carlo simulation

Consequences of Using Power Transforms as a Statistical Solution in Linear Mixed-Effects

Models of Chronometric Data

Models and methods examining basic cognition assign a central role to time, so much so that a common prerequisite for being a good model of cognition is to account for the time it takes to respond to a stimulus under various conditions. To be valid, analyses of chronometric data (e.g., response times [RTs] and fixation durations) must account for their consistent positive skew (Luce, 1986). This presents issues when analyzing RT data from individual trials using linear mixed-effects models (LMMs)—a class of statistical models whose use in cognitive psychology, particularly psycholinguistics, has grown over the last decade as part of a long-standing effort to develop ever better ways to generalize statistical tests across both participants and items (Clark, 1973; Raaijmakers, Schrijnemakers, & Gremmen, 1999; see Special Issue of the *Journal of Memory and Language* in 2008 on Emerging Data Analysis, specifically Baayen, Davidson, & Bates, 2008)—because in extending the standard linear model, LMMs also assume that the residuals, and by extension the outcome measure, are normally distributed. However, the positive skew of RTs entails that this assumption is typically violated in analyses of chronometric data.

Violating the normality assumption when fitting LMMs has been shown to produce unreliable standard errors and inference tests with potentially questionable results (Maas & Hox, 2004; Raudenbush & Bryk, 2002; Snijders & Bosker, 2012). Consequently, proponents of LMMs have advocated applying a power transform (e.g., log or inverse transform) to RT data prior to fitting LMMs to approximate residual normality, reduce positive skew, and reduce the impact of outliers (e.g., Baayen, 2008; Baayen & Milin, 2010), and this procedure has been adopted in many psycholinguistic studies that examine RT data. A survey of articles in *JML* from

January 2013 to December 2017 shows that this recommendation has been adopted in many studies, with power transforms applied in 31 (46%) out of the 68 articles where LMMs were fit to chronometric data¹. Accordingly, power transforms have been largely used as a “statistical solution” (Wagenmakers, Krypotos, Criss, & Iverson, 2012, p. 156) where data are transformed to better meet the assumptions of a data-analytic technique.

Unfortunately, the popularity with which power transforms are received as a statistical solution to nonnormality in LMMs of RTs is yet to be matched by a recognition of the complexities that power transforms bring to the analysis and interpretation of chronometric data. This mismatch is reflected in how researchers often overlook the idea that applying power transforms to the data not only reduces skew but also represents an implicit but strong claim about how the predictors of interest are hypothesized to relate to the observed RT measure. For example, applying a log transform to RTs implies that the researcher posits an exponential relationship between the predictors and RTs. Yet researchers who fit LMMs to transformed RTs still generally interpret the effects in the model as if the relationships are linear and that the predictors directly affect the observed RTs.

The popularity of power transforms as a statistical solution to non-normal data may have contributed to an impression that there is *one* appropriate scale of analysis—namely, whichever the scale is of the power transform that optimally normalizes the RTs. This idea implies that LMM analyses of non-optimally transformed/normalized RTs, raw RTs in particular, are strictly inappropriate because they violate the normality assumption. This position undermines not only a rich history of experimental research based on frameworks that directly infer the independence and interaction of information-processing stages from observed differences in raw RTs (Donders,

¹ The search was conducted in May 2018 using “Baayen, Davidson, & Bates (2008)” or “Baayen (2008)” or “lme4” and “reaction time” or “response time” or “fixation duration” as search terms on [sciencedirect.com/search/advanced](https://www.sciencedirect.com/search/advanced).

1969; Sternberg, 1969), but also the success of process models in showing how predictors can affect RTs in various scales where RTs approximate normality though not always optimally, such as the log scale (e.g., logarithmic activation growth model; Ulrich & Miller, 1993) or the inverse/rate scale (e.g., diffusion model; Ratcliff, 1978).

Moreover, to assume that the (optimally) normalized RTs are *the* appropriate measure of analysis can prevent researchers from acknowledging that differences in analyses performed on RTs in various scales can extend to issues beyond whether the normality assumption was met. Differences in results can arise simply because power transforms systematically alter units of duration: by compressing the long end of the RT distribution relative to the short end, power transforms downplay differences in long RTs. These changes may then reveal that predictors which interact in the transformed RT scale do not interact in the raw RT scale and vice versa. In fact, the possibility of obtaining different results from analyzing raw and transformed RTs due to issues unrelated to meeting the normality assumption is well-recognized in longstanding theoretical and statistical discussions about the use of power transforms (Loftus, 1978; Sternberg, 1969; Wagenmakers et al., 2012). However, only recently have these discussions resurfaced specifically in the context of using power transforms in LMMs of RTs (e.g., Balota et al., 2013; Kliegl, Masson, & Richter, 2010; Masson & Kliegl, 2013).

For example, word frequency and stimulus quality have been argued to have additive/independent effects on word recognition based on raw mean RT analyses (Borowsky & Besner, 1993; Yap & Balota, 2007; Yap, Balota, Tse, & Besner, 2008), which show equivalent increases in RTs to visually degraded words, as compared to visually intact words, for high- and low-frequency words. However, mean RT analyses disregard variability in RTs due to individuals dynamically adjusting their processing across an experiment (Kinoshita, Forster, &

Mozer, 2008; Masson & Kliegl, 2013). Masson and Kliegl (2013) argued that given RTs' sensitivity to properties of previous trials, interactions between word frequency and stimulus quality might manifest across trials, and mean RT analyses might be obscuring these interactions. Indeed, Masson and Kliegl found these interactions when they incorporated the previous trial's word frequency and stimulus quality into LMMs and inverse transformed their RTs to approximate residual normality. However, Balota et al. (2013) found that fitting LMMs to inverse-transformed RTs produced significant interactions between word frequency and stimulus quality across trials in two of three published datasets when those interactions were not obtained in the original ANOVAs nor in LMMs fit to raw RTs. While analyses on raw RTs suggest that the processing stages tapped by word frequency and stimulus quality are additive/independent (e.g., $650 \text{ ms} - 550 \text{ ms} = 100 \text{ ms}$ for high frequency words; $850 \text{ ms} - 750 \text{ ms} = 100 \text{ ms}$ for low-frequency words), inverse-transformed RTs could suggest otherwise (e.g., $(-1000/650) - (-1000/550) = 0.28$ for high frequency words; $(-1000/850) - (-1000/750) = 0.16$ for low frequency words).

When power transforms are performed as a statistical solution, differing results from analyses of raw and transformed RTs may give the impression that determining the “correct” scale of analysis is the primary issue at stake. But recognizing that power transforms re-express how predictors are hypothesized to relate to RTs offers another perspective: differing results serve as a valuable reminder that interactions that do not cross over are *scale dependent* in that they are removable (or borderline removable) by power transforms (Loftus, 1978; Wagenmakers et al., 2012). In this perspective, differences in results are not viewed as inconsistencies, nor as indicating that one of the models is “incorrect”, nor as showing that power transforms “distort” results in the raw RT scale. Rather, they simply reveal ways in which an interaction is scale-

dependent. Consequently, the concern shifts from identifying the “correct” scale of analysis to interpreting scale-dependent interactions with caution because their theoretical implications are not straightforward given that their presence depends on the scale of analysis. Indeed, Wagenmakers et al. (2012) argue that “in the absence of a specific process model, the safest conclusion is that interactions that do not cross over cannot be interpreted in terms of underlying psychological constructs” (p. 156). This admonition is warranted because researchers often do not specify process models yet provide theoretical interpretations for scale-dependent/non-crossover interactions—namely, that the predictors in the model directly affect the time it takes to perform the mental operations that culminate in the observed

Discussions of general constraints on interpreting scale-dependent interactions (Loftus, 1978; Wagenmakers et al. 2012), together with demonstrations that common transformations affect the presence of interactions in specific studies (Balota et al., 2013; Kliegl et al., 2010; Lo & Andrews, 2015; Masson & Kliegl, 2013), show that using power transforms in LMMs of RTs has consequences for the interpretation of results beyond preempting statistical issues from violating the normality assumption (Baayen, 2008; Baayen & Milin, 2010). Those consequences are particularly important when the focus is on the presence (or absence) of interactions between predictors. The literature survey discussed above show that this focus is common, with interactions in LMMs examined in 62 (91%) of the 68 *JML* articles. The common focus of psycholinguistic studies on interactions underscores the motivation to evaluate how the scale-dependence of effects in the population would affect statistical power to detect main effects and interactions in raw RTs and optimally transformed RTs. In the current project, we examined how often and in what ways we might expect to observe different results from LMMs fit to raw RTs and LMMs fit to transformed RTs due to scale dependence in effects.

Overview

We report two studies that sought to better understand how using power transforms in LMMs of chronometric data can generate different results due to scale-dependence in effects. RT data from two word-recognition megastudies—multisite efforts to develop comprehensive behavioral databases for broad sets of linguistic stimuli (see Keeulers & Balota, 2015 for an overview)—were analyzed, and Monte Carlo simulations were performed based on the LMMs fit to those megastudies. Study 1 was based on the *English Lexicon Project* (ELP; Balota et al., 2007) and Study 2 was based on the *Semantic Priming Project* (SPP; Hutchison et al., 2013)².

By performing simulations based on analyses of the megastudies, we demonstrated how often and in what ways LMMs fit to raw RTs and to transformed RTs could be expected to differ in datasets of typical sample sizes such as those analyzed by Balota et al. (2013). The principal advantage of analyzing RT data from megastudies is that the results are likely to be robust, more so than when smaller datasets are analyzed. The robustness of these results entails that the obtained estimates would make for more meaningful and realistic parameters to be used in the simulations than arbitrary parameters or estimates from small studies. Another advantage of analyzing megastudies is that their size and richness allow reasonably complex LMMs to be fit. Using these reasonably complex LMMs to generate the simulations ensures that the simulated RTs are a composite of different meaningful sources of variability. Analyzing RT data from two megastudies offers the additional advantage of making it possible to assess whether patterns of results are more consistent across the datasets and simulations when RT data are optimally

² This report is based on a project that consisted of three studies; two of those studies are reported here for brevity. The third study, based on the *Form Priming Project* (FPP; Adelman et al., 2014), is available in the Supplementary Materials (<https://osf.io/vyu65/>). The FPP contained 28 priming conditions, and we categorized 26 of those conditions as “orthographically related” and the other two as baseline conditions. While this dichotomy considerably simplified the analytic strategy in this third study, its results were similar to the two studies reported here.

transformed or left untransformed.

Both studies in this project consisted of three parts (see Figure 1 for a roadmap) that examined the consequences of power transforms for the main effects and interactions of two word-/item-level predictors (x_1 and x_2) and one person-level/individual difference predictor (z_1):

1. **Analysis of megastudy data.** Data from each megastudy were transformed using a log ($\ln(RT)$), an inverse square-root ($-1000/\sqrt{RT}$), and an inverse transform ($-1000/RT$), which in addition to the raw RT data resulted in four RT scales. The log, inverse square-root, and inverse transforms are special cases of the Box-Cox family of power transforms, which is generally expressed via the following piecewise function:

$$y^* = \begin{cases} \frac{y^\lambda - 1}{\lambda}, & \lambda \neq 0 \\ \ln(y), & \lambda = 0 \end{cases}$$

where the power parameter λ that would optimally normalize the data is estimated using a profile likelihood function (Box & Cox, 1964). Based on this function, the identity transform is approximated by $\lambda = 1$, and the farther λ is from 1, the stronger the transform required to optimally normalize the data. This implies that the inverse transform (approximated by $\lambda = -1$) is stronger than the inverse square-root transform (approximated by $\lambda = -0.5$) and that both of those transforms are stronger than the log transform. That is, distributions optimally normalized by the inverse transform will not be sufficiently normalized by the log transform, whereas distributions optimally normalized by a log transform will be “over-transformed” by the inverse transform so as to restore, if not increase, skew. A value of -1000 is multiplied to the inverse square-root and inverse transforms to magnify the transformed values and to maintain the rank-order of the RTs.

LMMs were fit to the data in each of the four scales. Afterwards, we used the Box-Cox procedure to estimate the λ that would optimally normalize the trial-level (i.e., cluster-specific) residuals and then created residual QQ plots³ for each of the LMMs to visually inspect that the power transform which produced the least skewed residuals corresponded to the λ identified by the Box-Cox procedure. Finally, we compared the fixed- and random-effects results across the raw and transformed LMMs to determine the implications of reducing RTs' positive skew using power transforms.

2. **Simulations based on raw model as generating process.** To assess the long-run performance of the raw and transformed LMMs, we performed simulations based on the LMMs fit to the megastudy data and compared how effectively the models estimated the input conditions that were specified in the simulations. Small-version datasets of the megastudy were generated based on the raw RT LMM in Part 1, and the subject and item sample sizes within those simulated datasets were varied. Afterwards, the same raw and transformed LMMs in Part 1 were fit to all the simulated data and then the performance of the models was assessed using several measures such as bias, coverage, convergence rates, statistical power, and false positive (Type I error) rates.

The raw LMM was used as the basis of the generating process of these simulations because the majority of cognitive theories describe processing as occurring over time and in doing so, they assume (implicitly or explicitly) that the relationship between observed RT and the latent cognitive factors measured by the predictors is reasonably linear (Lo &

³ QQ (quantile-quantile) plots indicate how the residuals are distributed relative to a normal distribution with mean 0 and variance σ^2 . These plots are created by arranging residuals in quantile order and pairing them with values of the normal distribution evaluated at those quantiles. A straight diagonal line would indicate normality and curvature or "snaking" would indicate nonnormality.

Andrews, 2015; but see Wagenmakers et al., 2012).

3. **Simulations based on optimally transformed model as generating process.** A second set of simulations was performed based on the LMM fit to the optimally normalized megastudy RTs, and the same performance and comparison measures in Part 2 were assessed. The purpose of these simulations is two-fold: first, they address the possibility that outcomes in Part 2 depended on using the raw LMM to simulate data. Second, they test whether there are advantages to fitting LMMs in the true scale of the generating process (e.g., would inverse LMMs perform better than raw LMMs when the inverse scale is the true scale of the functional relationship between the predictors and RTs?). In particular, these simulations test whether scale-dependent interactions would manifest most often in LMMs fit to the true (generating) scale.

Given this general approach, the current project's goals were to determine:

- **Whether different sets of RTs generated using the same task and obtained from similar populations are optimally normalized by the same power transform.** All RTs analyzed from the megastudies were obtained from young (college-aged) adults who performed a lexical decision task. Determining if LMM residuals are optimally normalized by the same power transform across the three studies would indicate whether optimally transforming RTs in LMMs leads to consistent rescaling of the relation between latent cognitive processes and observed data in studies that test participants from similar populations performing the same task.

- **The consequences for main effects of optimally transforming RTs.** Power transforms have been shown to increase statistical power for main effects in ANOVAs (Judd, McClelland, & Culhane, 1995; Levine & Dunlap, 1982; Ratcliff, 1993). Because LMMs are an extension of ANOVAs, we expect to replicate these results in LMMs.
- **The consequences for interactions of optimally transforming RTs.** LMMs have been advanced as a principled way of assessing individual and item differences in effects by allowing the simultaneous estimation of same-level and cross-level interactions. We determined whether optimal power transforms on RTs from very large datasets alter the results of statistical tests for these interactions. We also determined how power transforms affect statistical power for interactions in LMMs fit to studies of more typical sample sizes.
- **The consequences for random-effect patterns of optimally transforming RTs.** While fixed effects are the focus of most statistical inferences in LMMs, random effects reveal individual and item differences that are unaccounted in the specified model. Thus, random effects contain information that can motivate hypotheses about potential systematic individual and item variation in phenomena of interest. The current project determines how optimal transforms alter random-effect patterns and their interpretations.

Figure 1. Roadmap of procedures for each of the studies conducted in the current project.

Method

Study 1: English Lexicon Project (ELP)

The ELP (Balota et al., 2007) was the first word-recognition megastudy and remains the largest published descriptive database of word recognition, with 1,290 subjects responding to a subset of 40,481 words. We analyzed the lexical-decision subset of the ELP which consists of 816 subjects, each of whom responded to a ~1,700-word subset all the words tested in the ELP. Consequently, ~34 observations were made on each of the 40,481 words available in the ELP. More details about the procedure are available in Balota et al. (2007).

In Study 1, we focus on the following three predictors in the ELP and their interactions:

- 1) The target word's frequency (x_1), which is the target's $\log_{10}$ frequency per million words based on the HAL corpus (Lund & Burgess, 1996);
- 2) The target word's length in number of letters (x_2);
- 3) The subject's vocabulary (z_1), which was measured as the number of correct responses on a variant of the Shipley (1940) vocabulary test.

Study 1.1: Analyzing ELP Data

Data preprocessing. Only RTs from accurate responses to lexical targets (i.e., word strings correctly identified as words) were analyzed. Prior to screening the data for outliers, word frequency and word length were item grand-mean z-transformed, and vocabulary was person grand-mean z-transformed⁴. The data were then preprocessed as follows prior to fitting the LMMs: First, following Yap, Balota, Sibley, & Ratcliff (2012), observations that were 200 ms and shorter or 3000 ms and longer were excluded. Second, we applied the three power

⁴ Each unique item contributed one data point in the calculation of the grand means and z-scores for word frequency and length; each subject contributed one data point in the calculation of the grand mean and z-scores for vocabulary. The z-scores for continuous predictors in the subsequent studies were also calculated this way.

transforms (i.e., log, inverse square-root, and inverse) to the data and then identified outlying observations as those which were beyond ± 2.5 SDs of the subject's mean. Third, subjects' mean RTs for each scale were updated based on the remaining data, and subjects whose mean RTs were shorter than 400 ms or longer than 1000 ms (and their transformed-RT equivalents) were considered unusual and were dropped from further analyses. Because each RT scale produced different means and standard deviations, the percentage of observations dropped from further analyses due to the outlier criteria described above range from 4.86% to 10.41% of the total data across RT scales. Twenty-six subjects were further excluded due to missing vocabulary scores, resulting in 727 to 758 subjects being retained for analyses.

Fitting the LMMs and identifying the optimal power transform. The main effects of word frequency, word length, and vocabulary and all interaction effects were entered into cross-classified LMMs, one for each RT scale. We specified random intercepts for subjects and items and random subject slopes for the word frequency (x_1) and word length (x_2) effects. All LMMs (including those in subsequent studies) were fit with restricted maximum likelihood using the 1.1-12 version of the `lme4` package in R (Bates, Mächler, Bolker, & Walker, 2015; R Core Team, 2017)⁵.

After the models were fit, the optimal power transform for the preprocessed data was identified using the `boxcox` function in the 7.3-45 version of the `MASS` package in R (Venables & Ripley, 2002; R Core Team, 2017). The Box-Cox procedure revealed that the RTs are optimally normalized using $\lambda = -0.46$, which roughly corresponds to the inverse square-root transform. Trial-level residual QQ plots (Figure 2) showed that stronger power transforms

⁵ The general `lmer` formula for all the LMMs run in this project is $(RT \sim x_1 * x_2 * z_1 + (1 + x_1 + x_2 | Subject) + (1 | Item))$.

reduced residuals' positive skew more, but these reductions were offset by increases in negative skew.

Figure 2. Trial-level residual QQ plots of models fit to ELP data.

Because the four RT scales consist of different units of measurement, t -statistics were used as a standardized metric for comparing the models' fixed-effect estimates. Table 1 shows the t -statistics for each fixed effect across all four models. The t -statistics for x_1 and x_2 increased as stronger power transforms were applied to the data. However, all two-way interaction t -statistics systematically shrank as the transform became stronger, so much so that the t -statistic for the x_1z_1 effect no longer met the critical value of $|2|$ and is thus nonsignificant in the inverse model⁶. Remarkably, the positive three-way interaction estimate in the raw model reversed direction and became negative in all of the transformed models. That is, while the frequency-length interaction gets weaker as vocabulary increases in the raw model, the same interaction gets stronger as vocabulary increases in all the transformed models.

⁶ Because calculating degrees-of-freedom and p values in cross-classified LMMs is non-trivial, we followed the convention of setting the critical t value = $|2|$ as a reasonable approximation of significance given $\alpha = 0.05$ and given that the t -distribution turns into the standard normal distribution as sample sizes increase (see Baayen (2008) as well).

RT Scale	x_1 (Freq)	x_2 (Len)	z_1 (Voc)	x_1x_2 (Freq x Len)	x_1z_1 (Freq x Voc)	x_2z_1 (Len x Voc)	$x_1x_2z_1$ (F x L x V)
Raw	-73.49	45.19	-7.65	-22.38	6.89	-6.01	3.51
Log	-89.17	51.06	-8.21	-11.61	5.39	-5.44	-1.50
Inverse Sqrt	-94.47	54.39	-8.10	-7.45	3.55	-4.87	-3.09
Inverse	-95.97	57.14	-8.09	-3.61	1.12	-3.69	-5.11

Table 1. *t*-statistics of fixed effects from LMMs fit to ELP data. Effects that changed significance and direction from raw model are boldfaced. *Freq* = word frequency; *Len* = word length; *Voc* = vocabulary.

Beyond the fixed effects, the power transforms affected the magnitude of random-effect correlation patterns but not the direction of the patterns. Whereas the size of the word frequency and word length effects for each subject is strongly related to subjects' average RT in the raw scale ($r = -0.73$ and $r = 0.68$ respectively), these correlations manifest as only modest in the inverse square-root scale ($r = -0.33$ and $r = 0.43$ respectively; Table 2).

Model	Raw LMM			Inverse Square-Root LMM		
Fixed Effects	Estimate ($\hat{\gamma}$)	Std. Error	t	Estimate ($\hat{\gamma}$)	Std. Error	t
Intercept	737.30	4.03	183.00	-37.7298	0.1040	-362.75
x_1 (Frequency)	-57.78	0.79	-73.49	-1.4983	0.0156	-94.47
x_2 (Length)	43.89	0.97	45.19	1.0928	0.0201	54.39
z_1 (Vocabulary)	-31.58	4.13	-7.65	-0.8560	0.1056	-8.10
x_1x_2 (Freq x Len)	-7.91	0.35	-22.38	-0.0639	0.0086	-7.45
x_1z_1 (Freq x Voc)	5.02	0.73	6.89	0.0484	0.1366	3.55
x_2z_1 (Len x Voc)	-5.63	0.94	-6.01	-0.0906	0.0186	-4.87
$x_1x_2z_1$ (F x L x V)	0.70	0.20	3.51	-0.0137	0.0044	-3.09
Random Effects	SD	Correlations		SD	Correlations	
Subject						
Intercept	107.78	1		2.8284	1	
x_1	18.12	-0.73	1	0.3393	-0.33	1
x_2	23.78	0.68	-0.59	0.4798	0.43	-0.29
Item						
Intercept	62.74			1.5745		
Residual	191.60			4.4688		

Table 2. Parameter estimates from LMMs fit to raw and inverse-square-root-transformed ELP data. These estimates were used as the input parameters for the simulations performed in Studies 1.2 and 1.3 respectively.

Study 1.2: Simulations based on raw model as generating process

Generating process. The parameters used in the generating process of this simulation were integer-rounded values of the significant fixed-effect estimates (in this case, all effects) and of the random-effect estimates obtained from the raw model (summarized in *Estimate* and *SD* columns under *Raw LMM* section in Table 2). Because the raw model assumed normally distributed residuals, the residual variance estimate σ^2 cannot be directly used to generate trial-level residuals. Instead, to mimic the skew and variance observed in the raw megastudy RTs, a Gamma distribution with shape parameter = 0.1 and scale parameter = 3000 was used to generate trial-level residuals. Finally, an intercept parameter of 650 was added to the simulated values to approximate the intercept estimated by the raw RT LMM in Study 1.1.

Dataset generation and analysis. The datasets were generated to simulate a smaller version of the lexical-decision component of the ELP, where lexical-decision data to the same words are obtained from all subjects. Because word frequency (x_1) is -0.35 correlated with word length (x_2) in the ELP, word frequency and length values were generated from a standard bivariate normal distribution with a -0.35 correlation. Lastly, we generated vocabulary (z_1) values from a standard normal distribution based on the distribution of scores in the ELP.

The datasets were set up to have 20 to 200 subjects fully-crossed with 20 to 200 items in increments of 20, resulting in 100 sample-size conditions that more than encompass the typical sizes of psycholinguistic studies. In each dataset, a random subset of up to 5% of all the observations were deliberately excluded to simulate incorrect responses and missing data. All datasets were then analyzed using the same preprocessing and fitting procedure described in the Study 1.1.

A new dataset was generated whenever a warning or an error about convergence arose in

fitting the raw model. For each sample size condition, this procedure was repeated until there were 1,000 datasets where the raw model seamlessly converged.

Performance measures: Raw model. The consequences of violating the normality assumption by fitting LMMs to raw RT data were assessed by calculating percentage bias and coverage for each effect in the model. The effects of subject and item sample sizes on these performance measures were also assessed. Percentage bias was computed as $[(\text{estimate } \hat{\gamma} - \text{true value } \gamma) / \text{true value } \gamma] \times 100$, where positive values indicate that the parameter was overestimated and values between 5-10% are considered tolerable (e.g., Kaplan, 1989). For null/non-existing effects, bias was simply the estimate $\hat{\gamma}$. Coverage was determined as the percentage of occasions that the 95% confidence interval for estimate $\hat{\gamma}$ included the true value γ .⁷ Coverage provides a measure of whether the confidence intervals generated from the model tend to be conservative or permissive/anti-conservative, where percentages below 95% indicate conservative coverage.

Comparison measures: Transformed vs. raw model. Because the four RT scales consist of different units of measurement, our ability to compare the models fit to each RT scale was limited. Moreover, models that converge when fit to raw RT do not necessarily converge when fit to the transformed data and vice versa. Therefore, we first compared the models in terms of how consistent they were in converging. Subsequent comparisons were limited to datasets where both the raw model and its transformed counterpart both converged.

We then compared statistical power for the raw and transformed models in detecting effects present in the simulation and how much Type I error was incurred for null effects absent

⁷ The critical t value of $|2|$ was also used to compute the 95%-confidence intervals for all the parameter estimates.

in the simulation⁸, where power (Type I error) was defined as the proportion of models where the effect met the t critical value of $|2|$ when said effect was present (absent) in the simulation. To model how power and Type I error changed as a function of sample size, we fit loess curves to the power and Type I error estimates from all the sample size conditions and plotted those curves onto contour plots. To model how *differences* in power/Type I error between the transformed and raw models change due to sample size, we subtracted the estimates obtained from the raw and transformed LMMs at each sample size condition, fit loess curves to these difference estimates, and plotted these difference curves onto contour plots.

Since the use of power transforms in psycholinguistics has been primarily motivated by meeting the normality assumption, the main comparison of interest is between the raw model and the optimally transformed model based on the megastudy analysis (i.e., the inverse square-root model for the ELP data). Comparison measures for the other transformed models and more details about all the simulations are available in the Supplementary Material

(<https://osf.io/vyu65/>).

Results: Study 1.2

Performance measures: Raw model. Violating the normality assumption by fitting the raw model resulted in two consequences for the fixed effects. First, the raw model underestimated all the existing effects. Averaging across all subject and item sample sizes, the raw model underestimated the x_1 , x_2 , and z_1 main effects by 7.15%, 7.46%, and 9.88% respectively; the x_1x_2 , x_1z_1 , and x_2z_1 two-way interaction effects by 8.73%, 11.70%, and 12.15% respectively; and the $x_1x_2z_1$ three-way interaction by 10.08%. Moreover, changes in

⁸ In Studies 1.2 and 1.3, there was no effect for which Type I error was computed because all effects in the generating processes were non-zero.

sample size was largely unrelated to how much bias was incurred by the raw model.

Second, the raw model produced conservative 95%-confidence intervals for x_1 and x_2 at 91.3% and 92.8% respectively after averaging across all subject and item sample sizes, and increasing both subject and item sample sizes appeared to lower the coverage for these effects. On the other hand, the coverage for z_1 and all the interaction effects approximated the nominal 95% value and were less affected by changes in sample size.

Comparison measures: Transformed vs. raw model. The inverse square-root model converged in 92,976 of the 100,000 datasets where the raw model converged, and the datasets in which the inverse square-root model failed to converge tended to be those with fewer items. The results reported below were evaluated on the datasets where both models converged.

The contour plots in Figure 3a show predicted changes in power for the raw model as a function of subject and item sample sizes. As expected, power for all the effects increased as subject and item sample sizes increased. However, some effects were estimated to be more powerful than others: with 50 subjects and 50 items, x_1 and x_2 already had more than 0.80 in power, whereas power for z_1 and all the interaction effects were below 0.20 at these sample sizes. Similar patterns held for the inverse square-root model, except that power for the x_1x_2 interaction did not increase with increases in sample size as shown in Figure 3b.

Figure 3c shows the estimated differences in power between the raw and the inverse square-root model for each of the effects as a function of subject and item sample size. For the main effects, the two models differed in power only in small sample sizes, and this difference remained below 12%. However, for the x_1x_2 , x_1z_1 , and x_2z_1 interactions, the *raw* model was estimated to have as much as 37%, 62%, and 34% more power respectively than the inverse square-root model as sample sizes increased. This indicates that the raw model substantially

outperformed the inverse square-root model in detecting the interactions despite underestimating them. Lastly, the estimated difference for the three-way interaction remained below 4% across sample sizes.

Study 1.3: Simulations based on inverse square-root model as generating process

Generating process. The parameters used in the generating process of this simulation include all fixed- and random-effect estimates obtained from the inverse square-root model (summarized in *Estimate* and *SD* columns under *Inverse Square-Root LMM* section in Table 2). With the Box-Cox procedure revealing that the inverse square-root transform optimally normalized the RT data, a normal distribution with mean = 0 and variance = $4.4688^2 = 19.9706$ was used to generate trial-level residuals.

Dataset generation and analysis. This was identical to Study 1.2, except that the simulated inverse square-root RTs were back-transformed only into raw RTs to compare the performance of the models fit to these two RT scales.

Comparison measures: Inverse square-root vs. raw model. The same comparison measures were obtained as in Study 1.2. However, only the estimated differences in power between the raw and inverse square-root model are reported below because the patterns for power as a function of subject and item sample size were similar to those in Study 1.2.

Results: Study 1.3

The inverse square-root model converged in 92,661 of the 100,000 datasets where the raw model converged. As in the previous simulation, smaller item samples seemed to exacerbate this issue.

Differences between the raw and inverse square-root models' power estimates are shown in the power difference contour plots in Figure 3d. While differences in power estimates for the

main effects and the three-way interaction across all sample size conditions remained 7% or lower, the raw model's power for the x_1x_2 , x_1z_1 , and x_2z_1 interactions became as much as 16%, 48%, and 27% higher than the inverse square-root model's respectively as sample sizes increased. Again, the raw model outperformed the inverse-square root model despite the simulation being generated by the inverse-square-root LMM fit to the ELP data.

Overall, the results are similar to those observed in Study 1.2, thereby providing evidence that the choice of underlying generating process might not affect the influence of power transforms on results obtained from LMMs.

Discussion

Analysis of the ELP revealed that power transforms led to systematic changes in the t -statistics of the fixed effects. Stronger power transforms increased the t -statistics for the word-level main effects (x_1 and x_2), decreased those for the two-way interactions so much so that x_1z_1 was no longer significant in the inverse scale, and reversed the sign/direction of the three-way interaction. They also dampened random effect correlation patterns present in the raw scale.

These results were supported by the simulations. No differences in power were observed for the main effects between the raw and inverse square-root LMM presumably because the main effects were massive to begin with. But analogous to the decreases in the t -statistics for the two-way interactions as stronger power transforms were applied to the ELP data, the simulations revealed that the inverse square-root LMM was strongly outperformed by the raw LMM in detecting all two-way interactions as sample sizes increased. This result was obtained even when the generating process used in the simulation was the *inverse square-root LMM* fit to the ELP data, indicating that the advantage of the raw LMM did not depend on the generating process of the simulation.

Figure 3. (A) Power contour plots for the raw LMM as a function of subject and item sample size. (B) Power contour plots for the inverse square-root LMM as a function of subject and item sample size. (C & D) Power difference contour plots (inverse square-root – raw) as a function of subject and item sample size. Blue indicates that the raw LMM has more power than the inverse square-root LMM. Green indicates that the inverse square-root LMM has more power than the raw LMM. Figures 3a-3c were made from the simulations in Study 1.2 (LMM fit to the raw ELP data), whereas Figure 3d was made from the simulations in Study 1.3 (LMM fit to the inverse-square-root-transformed ELP data). The power contour plots for Study 1.3 show the same pattern as 3a and 3b (see Supplementary Material).

Study 2: Semantic Priming Project (SPP)

The SPP (Hutchison et al., 2013) is a megastudy of semantic priming in word recognition. In a typical semantic priming experiment, a target word (e.g., doctor) is preceded by a semantically related (e.g., nurse) or unrelated (e.g., bread) prime, and responses to target words are faster on average when preceded by a related prime (Meyer & Schvaneveldt, 1971). We analyzed the lexical-decision subset of the SPP which consists of data from 512 subjects, each of whom responded to half of 1,661 prime-target word pairs (i.e., 830 or 831) and filled out a battery of individual-difference measures. More details about the procedure are available in Hutchison et al. (2013).

In Study 2, we focused on the following three predictors in the SPP and their interactions:

- 1) The target word's frequency (x_1), which is the target's $\log_{10}$ frequency per million words in the SUBTLEX-US corpus (Brysbaert & New, 2009);
- 2) The prime's semantic relatedness (x_2), which indicates whether the target was preceded by a semantically related or unrelated prime, thus representing the semantic priming effect;
- 3) The subject's vocabulary (z_1), which is their average score for three vocabulary subtests from the Woodcock-Johnson reading battery (Woodcock, McGrew, & Mather, 2001).

Study 2.1: Analyzing SPP data

Data preprocessing. Only RTs from accurate responses to lexical targets (i.e., word strings correctly identified as words) were analyzed. Prior to screening the data for outliers, word frequency and vocabulary were item grand-mean and person grand-mean z-transformed respectively. The same preprocessing procedure in Study 1 was then applied to the SPP data,

except that instead of identifying outliers simply as observations that are ± 2.5 SDs from subjects' mean RTs, the non-iterative moving outlier criterion procedure recommended by Van Selst & Jolicoeur (1994) was performed for each RT scale because an experimental manipulation is included in the model. For every subject, this procedure adjusts the outlier criterion by the sample size observed for each experimental condition to account for the uncertainty introduced by small sample sizes. The maximum bounds for this criterion are ± 2.5 SDs from the mean for conditions that have sample sizes greater than 100 observations, and the criterion decreases with decreases in sample size. This preprocessing procedure excluded 2.82% to 5.17% of the total data across RT scales from further analyses. Five subjects and 15 target words were further excluded due to missing predictor values (word frequency, vocabulary), resulting in 1,646 items and 494 to 499 subjects being retained for analyses.

Fitting the LMMs and identifying the optimal power transform. The main effects of word frequency, semantic relatedness, and vocabulary and all interaction effects were entered into cross-classified LMMs, one for each RT scale. We specified random intercepts for subjects and items and random subject slopes for the word frequency (x_1) and relatedness (x_2) effects.

After the models were fit, the optimal power transform for the preprocessed data was identified as in Study 1. The Box-Cox procedure revealed that the RT data was optimally normalized using $\lambda = -0.95$ which roughly corresponds to the inverse transform. The trial-level residual QQ plots from the models supported the results of the Box-Cox procedure (Figure 4).

Figure 4. Trial-level residual QQ plots from models fit to SPP data.

The t -statistics across all four models show a milder pattern than that observed in the ELP (Table 3). For the main effects, the t -statistics from the transformed models were slightly larger than those from the raw model. For the two-way interactions, stronger power transforms resulted in greater changes to the t -statistics relative to the t -statistics in the raw model. Notably, the x_2z_1 interaction effect, whose t -statistic was distant from the critical value of $|2|$ in the raw scale, just surpassed this critical value and was significant in the inverse scale. Moreover, although still significant, the t -statistics of the other two-way interactions moved closer to the critical value as more powerful transforms were applied to the data. Lastly, transforming the data did not affect the three-way interaction's t -statistic.

RT Scale	x_1 (Freq)	x_2 (Relate)	z_1 (Voc)	x_1x_2 (Freq x Relate)	x_1z_1 (Freq x Voc)	x_2z_1 (Relate x Voc)	$x_1x_2z_1$ (F x R x V)
Raw	-31.95	-28.65	-4.12	13.43	8.91	-0.48	0.04
Log	-33.85	-32.31	-4.18	13.48	8.61	-1.28	0.43
Inverse Sqrt	-34.04	-32.89	-4.39	13.11	7.55	-1.76	0.84
Inverse	-33.60	-32.20	-4.40	12.24	6.71	-2.31	0.54

Table 3. t -statistics of fixed effects from LMMs fit to SPP data. Effects that changed significance from raw RT model are boldfaced. *Freq* = word frequency; *Relate* = semantic relatedness/priming effect; *Voc* = vocabulary.

Beyond the fixed effects, the power transforms also affected the estimation of random-effect correlation patterns. Whereas larger word frequency effects were associated with slower subjects in the raw scale ($r = -0.48$), the opposite relation appeared in the inverse scale ($r =$

0.36). Moreover, while the magnitude of the semantic priming effect was not related to subjects' mean RTs in the raw scale ($r = 0.06$), the effect was shown to be smaller for slower subjects in the inverse scale ($r = 0.48$) (Table 4).

Model	Raw LMM			Inverse LMM		
Fixed Effects	Estimate ($\hat{\gamma}$)	Std. Error	t	Estimate ($\hat{\gamma}$)	Std. Error	t
Intercept	664.01	5.01	132.47	-1.5975	0.0107	-149.09
x_1 (Frequency)	-29.73	0.93	-31.95	-0.0707	0.0021	-33.60
x_2 (Relatedness)	-21.26	0.74	-28.65	-0.0560	0.0017	-32.20
z_1 (Vocabulary)	-20.31	4.93	-4.12	-0.0464	0.0105	-4.40
x_1x_2 (Freq x Rel)	7.23	0.54	13.43	0.0132	0.0011	12.24
x_1z_1 (Freq x Voc)	5.00	0.56	8.91	0.0077	0.0011	6.71
x_2z_1 (Rel x Voc)	-0.35	0.74	-0.48	-0.0040	0.0017	-2.31
$x_1x_2z_1$ (F x R x V)	0.02	0.54	0.04	0.0006	0.0011	0.54
Random Effects	SD	Correlations		SD	Correlations	
Subject						
Intercept	109.86	1		0.2355	1	
x_1	9.16	-0.48	1	0.0190	0.36	1
x_2	11.50	0.06	0.14	0.0307	0.48	0.33
Item						
Intercept	30.04			0.0716		
Residual	163.07			0.3301		

Table 4. Parameter estimates from LMMs fit to raw and inverse-transformed SPP data. These estimates were used as the input parameters for the simulations performed in Studies 2.2 and 2.3 respectively.

Study 2.2: Simulations based on raw model as generating process

Generating process. The parameters used in the generating process of this simulation were integer-rounded values of the significant fixed-effect estimates (implying 0 for x_2z_1 and $x_1x_2z_1$) and of the random-effect estimates obtained from the raw model (summarized in *Estimate* and *SD* columns under *Raw LMM* section in Table 4). Setting up the generating process was similar to Study 1.2, except that a Gamma distribution with shape parameter = 0.15 and scale parameter = 1000 was used to generate trial-level residuals. Finally, an intercept parameter

of 600 was added to the simulated values to approximate the intercept estimated by the raw RT LMM in Study 2.1.

Dataset generation and analysis. The datasets were generated to simulate a typical semantic priming experiment. In this experiment, all subjects see the same target words, but half the targets are preceded by related primes. Prime relatedness (x_2) was counterbalanced in two lists so that targets preceded by related primes in one list are preceded by unrelated primes in the other list. Based on their distributions in the SPP, word frequency (x_1) and vocabulary (z_1) values were generated from standard normal distributions, and these values were sorted so that word frequencies are matched between the relatedness conditions.

The same data generation procedure was performed under the same subject and item sample size conditions as in Study 1 (ELP). However, the data were preprocessed as described in Study 2.1 (i.e., using the Van Selst & Jolicoeur (1994) procedure) to account for the prime relatedness manipulation (x_2) introduced in the simulations.

Performance and comparison measures. The same performance and comparison measures as in Study 1.2 were obtained.

Results: Study 2.2

Performance measures: Raw model. The consequences incurred for violating the normality assumption by fitting the raw model were similar to those in Study 1 (ELP). First, the raw model underestimated all the existing/non-zero effects. Averaging across all subject and item sample sizes, the raw model underestimated all the existing effects in the model: the x_1 , x_2 , and z_1 main effects by 5.99%, 0.34%, and 7.91% respectively; and the x_1x_2 and x_1z_1 interaction effects by 5.84% and 7.42% respectively. The average bias for the null effects x_2z_1 and $x_1x_2z_1$ were negligible. Changes in sample size were not related to how much bias is incurred by the

raw model.

Second, the raw model produced slightly conservative coverage for x_1 : averaging across all subject and item sample sizes, only 93.1% of the generated 95%-confidence intervals contained the true value of x_1 . Moreover, increasing both subject and item sample sizes appeared to lower the coverage for this effect. All other effects had coverages that approximated the nominal 95% value. For these effects, changes in sample size did not affect their coverage estimates.

Comparison measures: Transformed vs. raw model. The inverse model converged in 90,181 of the 100,000 datasets where the raw model converged, and as in Study 1 (ELP), the datasets in which the inverse model failed to converge tended to be those with fewer items.

The contour plots in Figures 5a and 5b show predicted changes in power/Type I error for the raw and inverse models as a function of subject and item sample size. As expected, power for all existing effects increased as subject and item sample sizes increased. Again, some effects were estimated to be more powerful than others: for instance, x_1 already had massive power at 0.80 with 50 subjects and 50 items, whereas x_2 's estimated power was only at least 0.50 and z_1 , x_1x_2 , and x_1z_1 's power estimates were below 0.20 with these sample sizes. Lastly, violating the normality assumption in the raw model did not increase Type I error rates for the null effects: regardless of subject and item sample size, Type I error estimates for x_2z_1 and $x_1x_2z_1$ remained below 7%.

Figure 5c shows the estimated differences in power/Type I error between the raw and the inverse model for each of the effects as a function of subject and item sample size. For x_1 , x_2 , and z_1 , the inverse model was estimated to have as much as 24%, 36%, and 9% more power than the raw model as sample sizes decreased respectively. As for interactions, the inverse model had

greater power over the raw model only in the x_1x_2 effect, where it had as much as 15% more power than the raw model as sample sizes increased. However, the raw model had greater power than the inverse model for the x_1z_1 effect, where it had as much as 12% more power than the inverse model as sample sizes increased. Differences in Type I error rates for x_2z_1 and $x_1x_2z_1$ between the raw and inverse model remained below 5% regardless of sample size.

Study 2.3: Simulations based on inverse model as generating process

Generating process. The parameters used in the generating process of this simulation were the significant fixed-effect estimates (implying 0 *only for* $x_1x_2z_1$) and random-effect estimates obtained from the inverse model (summarized in *Estimate* and *SD* columns under *Inverse LMM* section in Table 4). With the Box-Cox procedure revealing that the inverse transform optimally normalized the SPP data, a normal distribution with mean = 0 and variance = $0.3301^2 = 0.1090$ was used to generate trial-level residuals.

Dataset generation and analysis. This was identical to Study 2.2, except that the simulated inverse RTs were back-transformed only into raw RTs to compare the performance of the models fit to these two RT scales.

Comparison measures: Inverse vs. raw model. The same comparison measures were obtained as in Study 1.2. However, only the estimated differences in power between the raw and inverse model are reported below because the patterns for power as a function of subject and item sample size were similar to those in Study 2.2.

Results: Study 2.3

The inverse model converged in 88,379 of the 100,000 datasets where the raw model converged. As in the previous simulations, smaller item samples seemed to exacerbate this issue.

Differences between the raw and inverse models' power estimates are shown in the

power difference contour plots in Figure 5d. For the main effects, the inverse model's power for x_1 and x_2 was as much as 8% and 14% higher than the raw model's in smaller sample sizes respectively, and the difference between the two models remained below 5% for z_1 . For the interactions, the raw model's power for the x_1x_2 and x_1z_1 interactions became as much as 13% and 25% more than the inverse model's as sample sizes increased respectively, but the inverse model's power for the x_2z_1 interaction was as much as 10% more than the raw model's as sample sizes increased. Differences in Type I error for the three-way interaction between the two models remained below 5% across sample size conditions.

Some of the results were similar to those observed in Study 2.2: the inverse model had more power for the main effects at smaller sample sizes, the raw model had more power in detecting the x_1z_1 effect as sample sizes increased, and Type I error for the three-way-interaction remained at/below 5% regardless of sample size for both models. But some results were different: while the inverse model showed more power for the x_1x_2 effect when the generating process was based on the raw model, the raw model showed *more* power for the same effect when the generating process was based on the *inverse model*. Moreover, the inverse model showed greater power for the x_2z_1 effect in this simulation, but it did not exhibit greater Type I error for this effect in Study 2.2. This is likely due to the x_2z_1 effect being an existing/non-zero in this simulation, whereas it was a null effect in Study 2.2. Overall, these results are ambiguous about whether the choice of underlying generating process affects the influence of power transforms on LMMs, unlike Study 1 (ELP).

Discussion

Analysis of the SPP revealed that power transforms led to systematic changes in the t -statistics of the fixed effects, as in the analysis of the ELP. Stronger power transforms slightly

increased the t -statistics of the main effects, decreased those of x_1x_2 and x_1z_1 effects, and increased that of x_2z_1 to significance. They also substantially affected the estimation of random-effect correlation patterns.

These results were only partially supported by the simulations. Analogous to the slight increases in the t -statistics of the main effects as stronger power transforms were applied to the SPP data, the simulations revealed that the inverse LMMs generally had greater power than the raw LMM for main effects in small sample sizes. This is consistent with results from simulations performed on smaller sample sizes in the context of ANOVAs (Judd, McClelland, & Culhane, 1995; Levine & Dunlap, 1982; Ratcliff, 1993). Moreover, analogous to the systematic decrease in the t -statistic of the x_1z_1 interaction as stronger power transforms were applied to the SPP data, the inverse LMMs also became less sensitive in detecting the x_1z_1 interaction as sample sizes increased.

However, unlike the results in Study 1 (ELP), some results seemed to depend on the generating process of the simulation. In Study 2.2 where the LMM fit to the raw SPP data was the generating process, the inverse model was more sensitive in detecting the x_1x_2 interaction than the raw model as sample sizes increased. But in Study 2.3 where the LMM fit to the inverse-transformed SPP data was the generating process, the raw model surprisingly outperformed the inverse model in detecting the x_1x_2 interaction as sample sizes increased.

Figure 5. (A) Power contour plots for the raw LMM as a function of subject and item sample size. (B) Power contour plots for the inverse LMM as a function of subject and item sample size. (C & D) Power difference contour plots (inverse – raw) as a function of subject and item sample size. Blue indicates that the raw LMM has more power than the inverse LMM. Orange indicates that the inverse LMM has more power than the raw LMM. Figures 5a-5c were made from the simulations in Study 2.2 (LMM fit to the raw SPP data), whereas Figure 5d was made from the simulations in Study 2.3 (LMM fit to the inverse-transformed SPP data). The power contour plots for Study 2.3 show the same pattern as 5a and 5b (see Supplementary Material).

General Discussion

The influence of power transforms in LMMs fit to RTs was examined for lexical-decision data from two megastudies as shown in Figure 1. Part 1 consisted of analyses of raw and transformed data from each study while Parts 2 and 3 were simulations generated using parameter estimates obtained from the results of Part 1. The simulations in Part 2 used the raw (untransformed) model fit to the megastudy to generate data while those in Part 3 used the optimally transformed model to generate the data. For the most part, our summary of the critical results will not distinguish between these two types of simulations because their results were very similar and if anything, stronger results were obtained from the simulations where the optimally transformed model was used as the generating process. Table 5 provides a summary of how the analyses and simulations advanced the understanding of the consequences of using power transforms in LMMs of RT data.

1. Box-Cox analyses performed in conjunction with LMMs (Part 1) showed that power transforms reduced the positive skew of residuals for data for both megastudies. Simulations based on the raw model as the generating process (Part 2) showed that LMMs fit to the raw data were biased toward underestimating nearly all effects and tended toward conservative coverage for main effects. Together, these results show, as expected, that the power transforms improve model performance on these dimensions. However, the optimal power transform differed across the megastudies, with the Box-Cox λ ranging from -0.46 (approximately inverse square-root transform) for the ELP to -0.95 (approximately inverse transform) for the SPP.
2. The specified LMMs included two item-level main effects (x_1 and x_2) and a person-level main effect (z_1). For both megastudies, these main effects were statistically significant in

analyses of both raw and optimally-transformed data, and the t -statistics for z_1 negligibly changed between these two sets of analyses. However, the t -statistics for the item-level effects were greater in the optimally transformed LMMs than in the raw LMMs. The simulations supported this general pattern and showed either no difference in power for these effects between the raw and optimally transformed LMMs or greater power for the optimally transformed LMM in smaller sample sizes. This is consistent with earlier work on the effects of power transforms in ANOVAs (Judd et al., 1995; Levine & Dunlap, 1982; Ratcliff, 1993).

3. The specified LMMs included a single same-level interaction (x_1x_2 ; ELP: word frequency by length, SPP: word frequency by semantic relatedness) and several cross-level interactions (x_1z_1 , x_2z_1 , and $x_1x_2z_1$). While differences in the nature of the x_2 factor across the megastudies precludes the direct comparison of the megastudy analyses, consistent patterns were observed across the megastudies. For both megastudies, the same-level interaction and the x_1z_1 cross-level interaction (i.e., word frequency by subject vocabulary) were statistically significant in analyses of both raw and transformed data, but their t -statistics were smaller for transformed data than raw data. Further, the simulations showed that power for detecting these interactions was generally lower for transformed data relative to raw data and that the difference in relative power increased with increases in sample size.

In contrast, the results for x_2z_1 and $x_1x_2z_1$ cross-level interactions differed in some cases for analyses of raw and transformed data. The t -statistic for the $x_1x_2z_1$ interaction in the ELP went from a statistically significant positive value for the raw data to a statistically significant negative value for the transformed data, and the t -statistic for the x_2z_1 interaction in the SPP went from being well below the critical value for statistical significance for raw data to beyond the critical value for transformed data. Simulations for these additional cross-

level interactions showed that power was not consistently influenced by the transform, but Type I error rates for null interactions did not differ between analyses of transformed and raw data.

4. In both studies, the random-effects structure included the residual correlation between participants' RT intercepts and word-frequency (x_1) slopes. Analyses of raw data in both megastudies showed that after controlling for vocabulary, the effect of word frequency was still larger for slower participants. Transforming the data had inconsistent effects on these residual correlations across the megastudies. For the ELP, the magnitude of the correlation was reduced, but it still indicated a tendency for slower participants to show larger word frequency effects. In contrast, for the SPP, the direction of this correlation reversed such that slower participants tended to show smaller word frequency effects after controlling for vocabulary.

Results	Study 1 English Lexicon Project (ELP) Balota et al. (2007)	Study 2 Semantic Priming Project (SPP) Hutchison et al. (2013)
Part 1: Analysis of megastudy data		
Optimal (Box-Cox) power transform	• $\lambda = -0.46$, approximating the inverse square-root transform	• $\lambda = -0.95$, approximating the inverse transform
Influence of transform on main effects	• t -statistics increased for x_1 and x_2 ; negligible change for z_1	• t -statistics increased for x_1 and x_2 ; negligible change for z_1
Influence of transform on same-level (x_1x_2) interactions	• t -statistic for x_1x_2 decreased	• t -statistic for x_1x_2 decreased
Influence of transform on cross-level (x_1z_1 , x_2z_1 , $x_1x_2z_1$) interactions	• t -statistics for x_1z_1 and x_2z_1 decreased; $x_1x_2z_1$ reversed direction	• t -statistic for x_1z_1 decreased, x_2z_1 increased past critical value for significance
Influence of transform on random effects	• Reduction in correlation between subjects' RT intercepts and absolute magnitude of their word frequency slopes from $r = 0.73$ to $r = 0.33$	• Reversal of correlation between subjects' RT intercepts and absolute magnitude of their word frequency slopes from $r = 0.48$ to $r = -0.36$
Part 2: Simulations based on raw model as generating process		
Consequences of fitting LMMs to raw RTs for non-zero effects	• Main effects and interactions were underestimated; x_1 and x_2 had conservative coverage	• Main effects and interactions were underestimated; x_1 had conservative coverage
Relative power for detecting main effects in transformed compared to raw data	• Relative power for main effects was unaffected	• Relative power increased for all main effects in small sample sizes
Relative power for detecting same-level (x_1x_2) interactions in transformed compared to raw data	• Relative power decreased for x_1x_2 as sample sizes increased	• Relative power increased for x_1x_2 as sample sizes increased
Relative power/Type I error for detecting cross-level (x_1z_1 , x_2z_1 , $x_1x_2z_1$) interactions in transformed compared to raw data	• Relative power decreased substantially for x_1z_1 and x_2z_1 as sample sizes increased; unaffected	• Relative power decreased for x_1z_1 as sample sizes increased;

$x_1x_2z_1$) interactions in transformed compared to raw data	for $x_1x_2z_1$	Type I error unaffected for x_2z_1 and $x_1x_2z_1$
Part 3. Simulations based on transformed model as generating process		
Relative power for detecting main effects in transformed compared to raw data	• Relative power unaffected for main effects	• Relative power increased for x_1 and x_2 in small sample sizes; unaffected for z_1
Relative power for detecting same-level (x_1x_2) interactions in transformed compared to raw data	• Relative power decreased slightly for x_1x_2 as sample sizes increased	• Relative power decreased for x_1x_2 as sample sizes increased
Relative power/Type I error for detecting cross-level (x_1z_1 , x_2z_1 , $x_1x_2z_1$) interactions in transformed compared to raw data	• Relative power decreased substantially for x_1z_1 and x_2z_1 as sample sizes increased; unaffected for $x_1x_2z_1$	• Relative power decreased substantially for x_1z_1 as sample sizes increased and increased slightly for x_2z_1 ; Type I error unaffected for $x_1x_2z_1$

Table 5. Summary of results from the studies in the current project.

Recommendations

In psycholinguistics, power transforms have been used in LMMs of RTs largely as a statistical solution (e.g., Baayen, 2008; Baayen & Milin, 2010) rather than in conjunction with process models that map underlying cognitive processes to observed data. As noted, assumptions about that mapping are particularly important when evaluating scale-dependent interactions, which can be eliminated or created by monotonic nonlinear transforms. Here, we extend what has hitherto been mathematical (Loftus, 1978; Wagenmakers, et al. 2012) and case-study illustrations (Balota et al., 2013; Kliegl et al., 2010; Lo & Andrews, 2015; Masson & Kliegl, 2013) of scale-dependence: we simulated how frequently we might expect the scale-dependence of effects to influence results of RT analyses in studies of typical sizes, particularly in the context of LMMs. Our results indicate that researchers would maximize the detection of scale-dependent interactions and estimate random effects most consistently across studies by analyzing raw RTs

instead of optimally normalized RTs.

Accordingly, if (1) the interpretive framework mapping latent cognitive processes to RTs is implicit (i.e., broadly linear and in terms of direct “amount of time” required for processing) and (2) the goal is to detect scale-dependent interactions, we recommend that analyses be performed on raw RTs. Our recommendation is in response specifically to the popular use of power transforms as a statistical solution: because the benefits of using power transforms to reduce nonnormality are outweighed by concomitant theoretical and statistical considerations, researchers should have more compelling reasons beyond nonnormality to transform their data. Our recommendation is not intended to apply to analyses using explicit process models (e.g., drift diffusion model; Ratcliff, 1978). We further emphasize that our recommendation is not offered as any kind of solution to the theoretical problem of determining the “correct” scale in which to analyze RT. We elaborate on the considerations for recommending the analysis of raw RTs below.

Theoretical Considerations

Researchers generally interpret predictors as directly and linearly affecting observed RT measures. By analyzing raw RTs, they would be adhering to an analytic strategy that is already consistent with this interpretive framework. While there are issues with this framework (see Wagenmakers et al. (2012) for arguments against the interpretation of raw RT results), there is a rich history of robust experimental research based on raw RT results (Donders, 1969; Sternberg, 1969), suggesting that operating within this framework is not only intuitively appealing (i.e., cognitive processes occur over time, and observed RTs reflect distinct stages of these processes that are influenced by predictors; Sternberg, 1969) but also reliable. It is therefore unsurprising that researchers implicitly default to operating within this framework when analyzing

chronometric data. In fact, transforming RTs within this framework—particularly when it is done solely to address nonnormality in RTs—should be surprising because its theoretical implications are inconsistent with this framework.

Nevertheless, we recognize that there are valid theoretical motivations to analyze transformed RTs. Analyzing transformed RTs is appropriate when the nonlinear relationships between the predictors and observed RT implied by the power transform are supported by process models or prior evidence of nonlinearity. For example, comparing differences in effects between younger and older adults in raw RTs may be invalid because older adults' RTs are confounded by the general cognitive slowing associated with age (Salthouse, 1985). Likewise, comparing the effect of a manipulation over time in raw RTs may be inappropriate because of practice effects. By log transforming and inverse transforming RTs respectively, the effects of various factors on RTs would be modeled appropriately as a nonlinear function of age (Lo & Andrews, 2015) and of practice time.

However, predictors can exhibit unique nonlinear relationships with RTs—some predictors can be exponentially related to RTs whereas others can be inversely related to RTs. This implies that it is unlikely for there to be one “correct” scale that captures how all the predictors within a model relate to the observed RTs, much less how these predictors interact. Identifying the unique form of each predictor's relationship with RTs is not trivial, and specifying these relationships simultaneously in a standard LMM is quite complicated (i.e., while the model would still be linear in the *terms*, it would be nonlinear in the *parameters*). Moreover, even if process models and prior evidence suggest that several predictors have the same nonlinear relationship with RTs, they do not ensure that the nonlinear relationship specified in the model holds in the population (Wagenmakers et al., 2012).

These complexities highlight why our recommendation is not offered as a solution to determine a “correct” scale of analysis and why using power transforms as a statistical solution does not make researchers arrive at *the* correct scale of analysis. Ultimately, the decision to use a power transform should be driven by researchers’ goals and interpretive framework and should be accompanied by an acknowledgment of the statistical considerations involved in the (non-)use of the power transform.

Statistical Considerations

Sensitivity to interactions. Analyses of the megastudies (Part 1) showed that the test statistics for interactions that are significant across different RT scales were smaller for LMMs fit to optimally-transformed data than for LMMs fit to raw data. In the absence of an explicit process model, the consistency of these interactions’ statistical significance across all four RT scales in the megastudies (see Tables 1 and 3) suggest that these interactions hold across different ways of scaling the data in the population. However, the smaller interaction test statistics in optimally-transformed LMMs suggest that if we scale down the size of the megastudies to studies of more typical sample sizes, these interactions are more likely to be detected in raw RT LMMs. The simulations using both the raw generating model (Part 2) and the optimally-transformed generating model (Part 3) support this conclusion.

One possible reason for the greater sensitivity of raw LMMs to detect interactions comes from distributional analyses of chronometric data. Consider the cross-level interaction of word frequency and vocabulary, where individuals with higher vocabulary exhibit reduced word frequency effects. Results of the raw LMMs were most consistent with previous reports of this interaction (e.g., Butler & Hains, 1979; Chateau & Jared, 2000; Yap et al., 2012; Yap, Balota, & Tse, 2009): the *t*-statistic for this interaction is highest in raw LMMs across the megastudies and

raw LMMs generally had more power in detecting this interaction in the simulations than optimally transformed LMMs. Distributional analyses of RT data show that individual difference factors such as age, IQ, and working memory have their strongest effects on the long tail of the RT distribution (see Balota & Yap, 2011 for a summary). By compressing the long end relative to the short end of the RT distribution, power transforms could downplay the effects of individual differences, thereby reducing the power of transformed LMMs to detect interactions that involve individual differences. For the same reason, transformed LMMs may not have as much power as raw LMMs in detecting the interactions of predictors that affect different components of the RT distribution (i.e., modal/body vs. tail) more generally. This issue is compounded by the scale-dependence of interactions to influence the sensitivity and consistency of the LMM to detect the interaction.

Consistency of statistical outcomes. Scale-dependent interactions are challenging to interpret because the validity of inferences about these interactions depends on the validity of scaling assumptions that underlie them. When those assumptions are implicit, the challenge of interpreting scale-dependent interactions may be managed by evaluating the *consistency* of scale-dependent interactions.

Wagenmakers et al. (2012) suggest two criteria for evaluating this consistency. Their first consistency criterion is to treat an interaction as significant if the test statistic exceeds the critical value across a range of transformations (i.e., across-transform consistency), since that would show that the interaction occurs on an underlying level across some range of process models. Their second consistency criterion is to see which transformations, if any, show significant interactions across different subsets of the data (i.e., across-subset consistency).

Our results suggest an additional consistency criterion, one that considers the pattern

across studies (i.e., across-study consistency). Both megastudies consisted of data from young adults (college students) performing lexical-decision tasks to visually presented words, yet the optimal power transform differed between the two studies. Data generated by the same psychological processes should require the same power transform to link the underlying psychological constructs with observed behavior. Consequently, the selected power transform should lead to consistent results across similar datasets.

Assessing whether these three criteria have been met is easier in studies as large as the megastudies studied here because they provide substantial power to detect small effects and the observed patterns tend to be more reliable. Evaluating consistency becomes more difficult as sample sizes decrease because nonconvergence rates increase in smaller samples and consistency can only be evaluated when all the models being compared converge. Even when these models converge, it is not clear whether differences in results are due to reduction in power, less reliable results, and/or scale-dependence in the effects. Using power transforms to normalize the residuals would exacerbate this difficulty because the optimal transform is likely to differ across studies, even those that use highly similar tasks and participant samples as was the case for the three megastudies. A serious consequence of using different power transforms across similar studies is that the scales on which interactions are tested would be different, even if the underlying form of the relationship between the predictors and RTs is identical. This complicates the interpretation of patterns of results across studies. For example, the random-effect correlation of subjects' intercepts and word frequency slopes was consistent in magnitude and direction across the megastudies in analyses of the raw megastudy RTs, but this correlation differs across the megastudies in analyses of optimally transformed RTs. These differing patterns of random effects for the optimally transformed LMMs make it difficult to interpret the unexplained cross-

level interactions represented by these random effects.

Alternatives to power transforms in addressing nonnormality. The consequences of violating the normality assumption on statistical tests could be counteracted by estimating robust standard errors for the fixed effects (e.g., using the Huber-White sandwich estimator; see Raudenbush & Bryk, 2002; Snijders & Bosker, 2012). A more fundamental alternative is to forgo the normality assumption and acknowledge the positive skew of RT data by fitting generalized linear mixed-effect models (GLMMs) (Lo & Andrews, 2015). However, it is not clear that GLMMs are a practical solution to RT skew in psycholinguistic data. Maximum likelihood estimates cannot be analytically determined in GLMMs, making GLMMs more difficult to fit than standard LMMs and leading to higher rates of failure of the models to converge. These difficulties indicate that GLMMs might have to be drastically modified to a point where potentially invalid assumptions about empirical patterns would have to be made for the model to converge⁹ (e.g., stipulating that subjects' intercepts are unrelated to the magnitude of their word frequency effects; see Supplemental Material of Lo & Andrews, 2015). Thus, while GLMMs provide a direct way to address the positive skew in chronometric data, the extent to which this will be useful is not yet clear.

Conclusion

Applying power transforms to chronometric data involves statistical and theoretical considerations that override their value in addressing residual nonnormality in LMMs of RTs.

⁹ In the early stages of this project, we attempted to compare standard LMMs to GLMMs with an identity link and a Gamma distribution for the RTs (Lo & Andrews, 2015). However, despite the size of the megastudies, the structure we specified for the standard LMMs had to be drastically reduced (i.e., three main effects, random subject intercept, random subject slope for one item-level predictor) for the GLMMs to converge. Instead of simplifying the standard LMMs' structure to make them comparable to the GLMMs, we decided to retain the original structure of the standard LMMs and to focus on evaluating how scale dependence affects statistical power in LMMs of raw and transformed RTs. In doing so, we were able to fit models that were as complex as LMMs that psycholinguists tend to fit and set up generating processes that simulated reasonable approximations of real LDT data.

Inconsistency across studies in the type of power transform that optimally normalizes the residuals means that using power transforms as a statistical solution may greatly complicate the comparison of results across studies. This is because different studies—even those using very similar tasks and populations—would be optimally analyzed using different mappings of underlying cognitive processes to observed RT data. Whereas optimal power transforms increase test statistics and power to detect main effects in LMMs, they decrease test statistics and power to detect interactions which are typically the focus of research in psycholinguistics, while having negligible effects on Type I error rates. Optimal power transforms also alter random-effect patterns in ways that complicate interpretation of individual and item differences in fixed effects of interest. Together, these findings support the recommendation that researchers identify more compelling reasons beyond nonnormality to use power transforms when fitting LMMs to RTs. In the absence of these reasons, it is costly to use power transforms as a statistical solution in LMMs of chronometric data.

References

- Adelman, J. S., Johnson, R. L., McCormick, S. F., McKague, M., Kinoshita, S., Bowers, J. S., ... & Scaltritti, M. (2014). A behavioral database for masked form priming. *Behavior Research Methods*, 46(4), 1052-1067.
- Baayen, R. H. (2008). *Analyzing linguistic data: A practical introduction to statistics using R*. Cambridge University Press.
- Baayen, R. H., & Milin, P. (2010). Analyzing reaction times. *International Journal of Psychological Research*, 3(2), 12-28.
- Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59(4), 390-412.
- Balota, D. A., & Yap, M. J. (2011). Moving beyond the mean in studies of mental chronometry the power of response time distributional analyses. *Current Directions in Psychological Science*, 20(3), 160-166.
- Balota, D. A., Aschenbrenner, A. J., & Yap, M. J. (2013). Additive effects of word frequency and stimulus quality: The influence of trial history and data transformations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(5), 1563-1571.
- Balota, D. A., Yap, M. J., Hutchison, K. A., Cortese, M. J., Kessler, B., Loftis, B., ... & Treiman, R. (2007). The English lexicon project. *Behavior Research Methods*, 39(3), 445-459.
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255-278.
- Bates, D., Maechler, M., Bolker, B.M., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1-48.

- Borowsky, R., & Besner, D. (1993). Visual word recognition: a multistage activation model. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19(4), 813-840.
- Box, G. E., & Cox, D. R. (1964). An analysis of transformations. *Journal of the Royal Statistical Society. Series B (Methodological)*, 26(2), 211-252.
- Brysbaert, M., & New, B. (2009). Moving beyond Kučera and Francis: A critical evaluation of current word frequency norms and the introduction of a new and improved word frequency measure for American English. *Behavior Research Methods*, 41(4), 977-990.
- Butler, B., & Hains, S. (1979). Individual differences in word recognition latency. *Memory & Cognition*, 7(2), 68-76.
- Chateau, D., & Jared, D. (2000). Exposure to print and word recognition processes. *Memory & Cognition*, 28(1), 143-153.
- Clark, H. H. (1973). The language-as-fixed-effect fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior*, 12(4), 335-359.
- Donders, F.C. (1969). On the speed of mental processes. *Acta Psychologica*, 30, 412-431.
- Faust, M. E., Balota, D. A., Spieler, D. H., & Ferraro, F. R. (1999). Individual differences in information-processing rate and amount: implications for group differences in response latency. *Psychological Bulletin*, 125(6), 777-799.
- Hutchison, K. A., Balota, D. A., Neely, J. H., Cortese, M. J., Cohen-Shikora, E. R., Tse, C. S., ... & Buchanan, E. (2013). The semantic priming project. *Behavior Research Methods*, 45(4), 1099-1114.
- Judd, C.M., McClelland, G.H., & Culhane, S.E. (1995). Data analysis: Continuing issues in the

- everyday analysis of psychological data. *Annual Review of Psychology*, 46, 433-465.
- Kaplan, D. (1989). A study of the sampling variability and z-values of parameter estimates from misspecified structural equation models. *Multivariate Behavioral Research*, 24, 41-57.
- Keuleers, E., & Balota, D. A. (2015). Megastudies, crowdsourcing, and large datasets in psycholinguistics: An overview of recent developments. *The Quarterly Journal of Experimental Psychology*, 68(8), 1457-1468.
- Kinoshita, S., Forster, K. I., & Mozer, M. C. (2008). Unconscious cognition isn't that smart: Modulation of masked repetition priming effect in the word naming task. *Cognition*, 107, 623-649.
- Kliegl, R., Masson, M. E., & Richter, E. M. (2010). A linear mixed model analysis of masked repetition priming. *Visual Cognition*, 18(5), 655-681.
- Levine, D.W. & Dunlap W.P. (1982). Power of the F test with skewed data: Should one transform or not?. *Psychological Bulletin*, 92(1), 272-280.
- Lo, S., & Andrews, S. (2015). To transform or not to transform: using generalized linear mixed models to analyse reaction time data. *Frontiers in Psychology*, 6, 1-16.
- Loftus, G. R. (1978). On interpretation of interactions. *Memory & Cognition*, 6(3), 312-319.
- Luce, R. D. (1986). *Response times* (No. 8). Oxford University Press.
- Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instruments, & Computers*, 28(2), 203-208.
- Maas, C. J., & Hox, J. J. (2004). Robustness issues in multilevel regression analysis. *Statistica Neerlandica*, 58(2), 127-137.
- Masson, M. E., & Kliegl, R. (2013). Modulation of additive and interactive effects in lexical decision by trial history. *Journal of Experimental Psychology: Learning, Memory, and*

- Cognition*, 39(3), 898-914.
- Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, 90(2), 227.
- R Core Team (2017). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. Retrieved from <http://www.R-project.org/>.
- Raaijmakers, J. G., Schrijnemakers, J. M., & Gremmen, F. (1999). How to deal with “the language-as-fixed-effect fallacy”: Common misconceptions and alternative solutions. *Journal of Memory and Language*, 41(3), 416-426.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59-108.
- Ratcliff, R. (1993). Methods for dealing with reaction time outliers. *Psychological Bulletin*, 114, 510-532.
- Raudenbush, S. W., & Bryk, A. S. (2002). *Hierarchical linear models: Applications and data analysis methods* (Vol. 22nd ed.). Sage.
- Salthouse, T. A. (1996). The processing-speed theory of adult age differences in cognition. *Psychological Review*, 103(3), 403-428.
- Schmider, E., Ziegler, M., Danay, E., Beyer, L., & Bühner, M. (2010). Is it really robust? Reinvestigating the robustness of ANOVA against violations of the normal distribution assumption. *Methodology: European Journal of Research Methods for the Behavioral & Social Sciences*, 6(4), 147-151.
- Seidenberg, M. S. (1985). The time course of phonological code activation in two writing systems. *Cognition*, 19(1), 1-30.
- Selst, M. V., & Jolicoeur, P. (1994). A solution to the effect of sample size on outlier

- elimination. *The Quarterly Journal of Experimental Psychology*, 47(3), 631-650.
- Shipley, W. C. (1940). A self-administering scale for measuring intellectual impairment and deterioration. *Journal of Psychology*, 9, 371-377.
- Smith, N. J., & Levy, R. (2013). The effect of word predictability on reading time is logarithmic. *Cognition*, 128(3), 302-319.
- Snijders, T. A. B., & Bosker, R. J. (2012). *Multilevel analysis: An introduction to basic and applied multilevel analysis* (Vol. 22nd ed.). Sage.
- Sternberg, S. (1969). The discovery of processing stages: Extensions of Donders' method. *Acta Psychologica*, 30, 276-315.
- Venables, W. N. & Ripley, B. D. (2002). *Modern applied statistics with S* (4th ed.). NY: Springer.
- Wagenmakers, E.-J., Krypotos, A. M., Criss, A. H., & Iverson, G. (2012). On the interpretation of removable interactions: A survey of the field 33 years after Loftus. *Memory & Cognition*, 40(2), 145-160.
- Wagenmakers, E.-J., Ratcliff, R., Gomez, P., & McKoon, G. (2008). A diffusion model account of criterion shifts in the lexical decision task. *Journal of Memory and Language*, 58, 140-159.
- Woodcock, R. W., McGrew, K. S., & Mather, N. (2001). *Woodcock-Johnson tests of achievement*. Itasca, IL: Riverside Publishing.
- Yap, M. J., & Balota, D. A. (2007). Additive and interactive effects on response time distributions in visual word recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33(2), 274-296.
- Yap, M. J., Balota, D. A., Sibley, D. E., & Ratcliff, R. (2012). Individual differences in visual

word recognition: Insights from the English Lexicon Project. *Journal of Experimental Psychology: Human Perception and Performance*, 38(1), 53-79.

Yap, M. J., Balota, D. A., Tse, C. S., & Besner, D. (2008). On the additive effects of stimulus quality and word frequency in lexical decision: evidence for opposing interactive influences revealed by RT distributional analyses. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34(3), 495-513.

Yap, M. J., Tse, C. S., & Balota, D. A. (2009). Individual differences in the joint effects of semantic priming and word frequency revealed by RT distributional analyses: The role of lexical integrity. *Journal of Memory and Language*, 61(3), 303-325.