Are Reaction Time Transformations Really Beneficial?

Pele Schramm¹ & Jeffrey N. Rouder¹

¹ University of California, Irvine

Author Note

Correspondence concerning this article should be addressed to Pele Schramm, Social and Behavioral Science Gateway, Irvine, CA. E-mail: pschramm@uci.edu

Abstract

We investigate whether or not the common practice of transforming response times prior to conventional analyses of central tendency yields any notable benefits. We generate data from a realistic single-bound drift diffusion model with parameters informed by several different typical experiments in cognition. We then examine the effects of log and reciprocal transformation on expected effect size, statistical power, and Type I error rates for conventional two-sample t -tests. One of the key elements of our setup is that RTs have a lower bound, called the shift, which is well above 0. We closely examine the effect that different shifts have for the analyses. We conclude that logarithm and reciprocal transformation offer no gain in power or Type I error control. In some typical cases, reciprocal transformations are detrimental as they lead to a lowering of power.

Keywords: Response Times, Data Analysis

Are Reaction Time Transformations Really Beneficial?

It is common in experimental psychology to use response times as a dependent variable that indexes overall performance. Famous examples include Sternberg’s demonstration of exhaustive short-term memory scanning (Sternberg, 1966) and Shepard and Metler’s demonstration of mental rotation (Shepard & Metzler, 1971). Response times are ubiquitous for assessing the effect of context on a target, and examples include classic Stroop and Simon effects (Simon, 1969; Stroop, 1935) as well as the more recent variants such as the weapons-priming tasks (Correll, Park, Judd, & Wittenbrink, 2002). Overall, the analysis of response times in high-accuracy experiments has been a time-tested, popular, and fruitful approach in experimental psychology.

The question then is how to analyze such data. In this paper, we consider testing whether response time is affected by a covariate. A common example is the Stroop effect. Here, participants identify the display color of a color term (e.g., the word *RED*). The display color is either the same as the color term (e.g., the word *RED* displayed in red) or different from the color term (e.g., the word *GREEN* displayed in red). Incongruent words slow response times, and accurately detecting this effect is a common goal.

The vast majority of researchers in these cases use familiar, well-trodden linear-model statistical tests such as *t*-tests, regression, and ANOVA, again with the goal of stating whether covariates affect RT or not. What we have noticed in the literature, however, is that while most researchers do so on untransformed variables, there are many examples of using transformations before analysis, most commonly either log or reciprocal transformations. Some examples of the use of a logarithmic transformation can be found in Harmon-Jones and Gable (2009) and Levav and Fitzsimons (2006); examples of reciprocal transformation can be found in Bem (2011) and Massar, Lim, Sasmita, and Chee (2016). Is transforming wise? Indeed, the genesis of this paper came from a friendly lunchtime disagreement among

the authors about the wisdom of transformations. Over burritos, the first author expressed his discomfort about using normally distributed models for RTs, a recent approach favored by the second author (Haaf & Rouder, 2017, 2019; Rouder & Haaf, in press; Thiele, Haaf, & Rouder, 2017). The first author recommended placing normal models on the logarithm of response time. This paper provides an assessment of whether it is wiser to perform linear-model tests on transformed or untransformed response time.

The initial recommendation to transform parameters comes from the seminal work of Box and Cox (1964). These authors show that with simple, univariate transforms, skewed data may be brought better in line with the normal distributional assumptions. As a result, statistical techniques that assume normally distributed residuals may be used with less worry about the effects of violating distributional assumptions. Based on this advice, some researchers recommend transformations. For example Wilcox (1998) makes an appeal for robust statistics including using transforms. Whelan (2008) reviews the benefits of transformation, which for RT include minimizing the impact of slow-response outliers while maintaining good power, as well as a concern over the interpretation of transformed data. Skeptics of transformed variables include Ratcliff (1993), who provides a small simulation study, and Lo and Andrews (2015), who recommends generalized linear mixed models as an alternative on the basis of a case study in word naming.

In this paper, we seek to go beyond the previous assessments. Our approach relies on both an analytic assessment of the noncentrality parameter that determines the power in statistical tests as well as a simulation check among a broad array of realistic conditions. Here, we pay particularly close attention to generate as realistic representations of RTs that we know. We opt for a shifted, one-bound diffusion model of perception (Laming, 1968; Ratcliff, 1978; Stone, 1960) as a generative model. The parameters we use in the generative model are directly informed by trends in the empirical literature, and consequently, our results are directly applicable in real world situations.

Rationale For Transforms

To understand the dilemma facing researchers, it is helpful to review the following properties of response times. First, response times are unimodal with a skewed upper tail (Luce, 1986; Van Zandt, 2000). Second, manipulations that slow the distribution tend to increase the mean and standard deviation together with relatively small effects in higher moments such as skew. (Rouder, Pratte, & Morey, 2010; Wagenmakers & Brown, 2007). Third, effects across conditions may be difficult to detect because there is excessive trial-by-trial variability as well as much variability across people (Haaf & Rouder, 2017). Figure 1, Panel A, provides an example of these properties. Consider the two distributions that are denoted with solid lines. These two are for two experimental conditions such as *congruent* and *incongruent*. The distributions are unimodal with a large degree of skew for the upper tail. The distributions are highly similar indicating that the effect is small relative to the trial-by-trial variability. Finally, the difference between conditions is manifest in both the mean and variance.

These properties taken together indicate that a transform that disproportionately affects the long tail may be helpful in reducing skewness and stabilizing variance. Indeed, the aforementioned logarithm and reciprocal transform disproportionately reduce the long tail. Therefore, it seems reasonable that these transforms are helpful in assessing small effects. Moreover, both logarithmic and reciprocal transform have possible process interpretation. Ulrich and Miller (1993) provide four separate process scenarios that generate lognormally distributed data with most centered around multiplicative gain and attenuation from attentional filters. Likewise, Reddi and Carpenter (2003) describe a ballistic accumulator model that naturally gives rise to response times which are the reciprocal of a truncated normal distribution. So, given this constellation of facts, it certainly seems plausible that RT transformations are beneficial in analysis.

The main goal was to see how easily we could detect an effect between two conditions. Hence, the main characteristics we assess is the performance of a conventional equal-variance t -test. For each transform, we ask which has the highest power to detect an effect with a Type I error rate maintained at the specified level. We can get a rough guideline to the first question, about power, by studying the *effect size* between two different distributions. The effect size is the difference in expected values scaled by a measure of averaged standard deviation.¹ It is a known constant for any two distributions, for example, the effect size between the solid distributions in the untransformed case is 0.34. The effect size, for fixed N directly determines the expectation of the *noncentrality parameter* of the noncentral-T, which is monotonically related to power. Hence, bigger effect sizes correspond to more power.

How does transformation affect the effect sizes between the distributions? The effect sizes for the logarithmic and reciprocal transformed distributions are, respectively 0.33 and 0.28. Notice that these values are lower than that for the untransformed distribution. This simple effect-size analysis provides our first clue—transforming data will not increase power, and, at least for the reciprocal transform, it may even lower it. Unfortunately, while studying the effect size provides strong guidance, it is not the final answer to our question. First, it does not address the level of the test, that is, the real Type I error rate. Second, while perfectly predictive for power in the asymptotic limit, the effect-size analysis may not hold with finite samples. Nonetheless, we should be quite skeptical of any gains from transformations off the bat.

Those Pesky Shifts

There is one additional property of response times that has been overlooked by those using transforms. This property is that all response time distributions have a substantial shift away from zero. Shifted distributions are shown in Figure 1 as dashed lines, and these

¹ Let X and Y be two distributions. The effect size is $(E(X) - E(Y))/\sqrt{.5(V(X) + V(Y))}$.

particular distributions have a shift value of .5 sec. Notably, there is no possibility of a response below this .5 sec value. Shifts in RT distributions are ubiquitous (Luce, 1986; Rouder, 2005); indeed, we know of no data set with mass near zero. Figure 2A, from Rouder, Province, Morey, Gomez, and Heathcote (2015), provides an example. The data here come from 94 participants in lexical decision experiment from Gomez, Ratcliff, and Perea (2008) where people responded in about 700ms on average. The plotted values are the shift parameter from a lognormal distributional fit as performed in Rouder et al. (2015). The shift represents the point below which there is no mass. The thin lines are 95% credible intervals, and shifts are located far from zero for all participants.

Consideration of shifts, in our opinion, simply reflects the nature of response times themselves. After all, it is a reality of neurophysiology that the mere perception of an unanticipated stimulus is doomed to be delayed post exposure at least by the amount of time it takes for the chains of action potentials from the retina to relevant brain areas to run their course. Likewise, the same can be said for the chain of action potentials required to translate the decision into the proper button press. No wonder so many RT modeling efforts include shift parameters (e.g., Dzhafarov, 1992; Ratcliff, 1979; Usher & McClelland, 2001; Vickers & Smith, 1985).

The presence of these shifts pose a serious challenge to the rationale for transforms. Figure 1, Panels B and C shows the effects of shifts. Here, the shift can dramatically change the shape and scale of the resulting transformed distribution. For example, unshifted log-transformed distributions (solid lines) are more symmetric than shifted log-transform distributions (dashed lines). Likewise, shifted reciprocal transformed distributions (dashed lines) have negative skew rather than being symmetric. Moreover, not only is the variance not stabilized, the degree of variance is a function of the shifts, indicating that the transforms may work in some paradigms better than others.

A Simulation Approach

To understand the effects of transforming response time in realistic situations, we decided to use a simulation method. The results of a simulation are only as good as the choices made for inputs. We strove for highly realistic choices that characterized common research. We decided to use the diffusion model of perception as our base generative model. The key assumption in diffusion models is that moment-to-moment evidence accumulates gradually until there is enough total evidence to support a response. For tasks where accuracy is relatively high, say those commonly used for examining working memory, response conflict, and perceptual representations, the one-bound diffusion model is highly appropriate (Rouder et al., 2010). A visualization of the evidence-accumulation process is shown in Figure 2B.

The distribution of absorption times of the one-bound diffusion process is known as an inverse Gaussian (Chhikara & Folks, 1989) or a Wald distribution (Stone, 1960). When parameterized with drift-rate (ν), bound(λ), and shift(ψ), the density may be expressed in closed-form as:

$$f(t; \psi, \nu, \lambda) = \frac{\lambda}{\sqrt{2\pi}} (t - \psi)^{-3/2} \exp \left(- \frac{[\lambda - \nu(t - \psi)]^2}{2(t - \psi)} \right)$$

Where $t > \psi > 0$, $\lambda > 0$, $\nu > 0$. The close form nature of this density makes it far more convenient than the two-bound diffusion model (the density is given by the convergence of an infinite sum). Moreover, sampling from the density may be performed by transforming chi-squared deviates as described in Michael, Schucany, and Haas (1976) and as implemented in the STATMOD package in R.

The inverse Gaussian may be parameterized in terms of shift, drift rate, and bound, as above, or, alternatively, in terms of shift, scale, and shape. For guidance about which parameterization is best for empirical phenomena, we rely on Rouder et al. (2010) who

competitively tested the two different parameterizations. In one model, stimulus strength affected drift rate and bound was held constant; in the other model stimulus strength affected scale and the shape was held constant. The results were unequivocal—the constant bound model was superior. Over 90,000 observations, RT distributions showed small but detectable shape differences remarkably like those in Figure 1A. Given Rouder et al’s demonstration we manifest the difference between our simulated conditions in drift rate.

We start with three modal response-time experiments. The first is what we term the *perceptual experiment*, and the defining characteristic is fast overall response times and small differences between conditions. We performed several variants of the fast theme, and one of these, the most realistic, serves as our *canonical version*. The canonical version of this experiment has an overall mean of 0.30 sec with a 0.02 sec difference between conditions. Table 1 shows the corresponding settings of shift, bound and drift rates for power as well as a common drift rate for level. The canonical version has 60 trials per condition, and the entry is shown in the top row of Table 2. The remaining 8 rows show different versions made by varying the shift and the total number of trials per condition. We also simulated data from hypothetical *attention* and *linguistic* experiments. The settings for the attention experiments are designed to capture mid-level cognition experiments which take about 6/10ths of a second and have effects on the order of 50 ms. The particular settings are shown in Table 1, and the versions populate the middle set of rows in Table 2. Likewise, the linguistic set of experiments is designed to capture above-second tasks that typically have larger effects.² In addition to simulating data from these hypothetical experiments, we report the effect size between the distributions in each case for each transformation (Table 3).

² The number of trials in the linguistic experiments is chosen to be greater than those in the perception and attention experiments so that the power is roughly comparable.

Simulation Results

The results for power and level for each of the versions of each of the experiments are shown in Table 2. Under *Power* and *Level*, there are three columns each. The first, *RT* denotes results for untransformed distributions; the remaining two are for the logarithm and reciprocal transforms, respectively. The main results here are easy to characterize. For most cases, it matters little in power or level whether response times are transformed or not, though no transformation and log transformations seem to do marginally better than reciprocal transformations in most instances. From 3 we can also see this reflected in effect size. Although the *t*-test appears to be fairly robust in most of these cases under transformation, there is no gain that we can see in transforming the data.

To understand the full range of effects of transformation on *t*-test results, we plotted the receiver operating characteristic (ROC). The left column of Figure 3 shows the case for the canonical versions of each modal experiment. The *x*-axis of each ROC is the nominal level for a test, or, more conventionally, the α setting. The *y*-axis is the rejection rate, which describes both the power of a test (when there are true differences between conditions) and the level of a test (when there are no differences between conditions). These ROCs show that in all cases, the level is fairly close to the nominal value, which is heartening. Moreover, the overlap of all curves in the perception and attention experiments show how robust the *t*-test is to the transformations for these ranges of observations. The outlier is the linguistics experiments, but here, contrary to conventional wisdom, the highest power is attained with the untransformed distributions. Therefore, we conclude that there is nothing to gain by transforming the data.

Variation in Shift

One feature of the above simulation is that the parameter values were held constant across all replications of a simulation. We therefore call these the fixed parameter simulations. It is common, however, to add variability to the parameters of the diffusion model in application. For example, Ratcliff, Thapar, and McKoon (2001) added variability to the drift rate, bound, and shift to more faithfully model their data. Schwarz (2001) added exponential variability in the shift, but we believe that adding a small-degree of normal variation is more realistic. As shown in 4, this addition has the effect of making the decrease of the left tail more gentle.

We ran the simulations with a small degree of variability in the shift parameter. The shift used to generate each simulated reaction time was drawn from a normal, centered on the shift values indicated in Table 5, and with a standard deviation of 15% of the mean shift value. Power and level results are shown in Table 5, effect sizes in 4 and ROCs are shown in the right column of Figure 3. In the tables, there are no entries for zero shift in this case because we did not wish to entertain negative shift values.

In each of the cases effect sizes appear marginally better untransformed, closely followed by the log transform, and then by the reciprocal transform. This pattern is also for the most part reflected in our power simulations. Again, the results of the level simulations seem to be more or less as they should be in all cases. The ROC curves demonstrate more of the same, though with a random shift they seem to favor untransformed data ever so slightly more.

Discussion

In this paper, we consider the effects of transforming response times in realistic cases through analysis of the effect sizes of distributions and through simulation. Our results are

clear. Transforming variables offers no more power nor any better level control than not transforming. In specific cases, certain transforms, notably the reciprocal transform, provided for lower power. Thus, we see no reason to transform response times in establishing effects.

The remaining question then is about the generality of the results. We think they hold fairly broadly when the goal of the researcher is to establish effects through linear model analysis. For example, we would expect no increased power from transforming if our goal was to detect whether RT covaried with variables such as word frequency. Likewise, we see no reason to transform variables in ANOVA analyses where the goal is to test main effects and interactions. Simply put, if the goal is to establish the presence or absence of effects, then transformation holds no advantages that we can identify.

A corollary to this result is that normal models may be used for response time when the goal is to establish nominal or ordinal relationships. Consider the goals of Haaf and Rouder (2017) who asked what they call the “does everybody” question. We know that *on average* people respond more quickly to congruent items than incongruent ones in the Stroop task. Does this ordering, congruent-faster-than-incongruent, hold for all people, or, alternatively, are there people who truly have no Stroop effect or have a pathology where they truly respond more quickly to incongruent items? Given that their questions are about detecting the presence and direction of effects, the results here indicate that the normal is not problematic. Indeed, Thiele et al. (2017) use simulations to show normal models of response time distributions are perfectly reasonable for such ordinal questions.

Researchers interested in modeling the functional relationship between RT and covariates—say whether RT goes as a power function or an exponential function of a covariate—do need more accurate models of RTs. Rouder, Tuerlinckx, Speckman, Lu, and Gomez (2008) showed that a main concern is modeling the shifts. A failure to do when shifts are present leads to systematic distortions of the functional relationship between RT and covariates. The presence of shifts that vary across people is a strong challenge because

models in the linear mixed model and generalized linear mixed model family are unable to capture individual-specific shifts. Fortunately, researchers do have an option. Bayesian hierarchical models are perfectly suited for the challenge, and recent developments in general purpose software such as `stan` (Carpenter et al., 2017) and JAGS (Plummer, 2003) make analysis with these models convenient for nonexpert researchers. There are now several excellent sources for guidance on hierarchical models in cognition including Kruschke (2014) and Lee and Wagenmakers (2013).

References

- Bem, D. J. (2011). Feeling the future: Experimental evidence for anomalous retroactive influences on cognition and affect. *Journal of Personality and Social Psychology*, *100*, 407–425. Retrieved from <http://dx.doi.org/10.1037/a0021524>
- Box, G. E., & Cox, D. R. (1964). An analysis of transformations. *Journal of the Royal Statistical Society: Series B (Methodological)*, *26*(2), 211–243.
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Bettencourt, M., . . . Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, *76*.
- Chhikara, R. S., & Folks, J. L. (1989). *The inverse Gaussian distribution: Theory, methodology, and applications*. New York: Marcel Dekkar.
- Correll, J., Park, B., Judd, C., & Wittenbrink, C. (2002). The police officer's dilemma: Using ethnicity to disambiguate potentially threatening individuals. *Journal of Personality and Social Psychology*, *83*(6), 1314–1329.
- Dzhafarov, E. N. (1992). The structure of simple reaction time to step-function signals. *Journal of Mathematical Psychology*, *36*, 235–268.
- Gomez, P., Ratcliff, R., & Perea, M. (2008). The overlap model: A model of letter position coding. *Psychological Review*, *115*(3), 577.
- Haaf, J. M., & Rouder, J. N. (2017). Developing constraint in Bayesian mixed models. *Psychological Methods*, *22*(4), 779–798.
- Haaf, J. M., & Rouder, J. N. (2019). Some do and some don't? Accounting for variability of individual difference structures. *Psychonomic Bulletin and Review*. Retrieved

from <https://psyarxiv.com/zwjtp/>

Harmon-Jones, E., & Gable, P. A. (2009). Neural activity underlying the effect of approach-motivated positive affect on narrowed attention. *Psychological Science*, 20(4), 406–409.

Kruschke, J. K. (2014). *Doing bayesian data analysis, 2nd edition: A tutorial with r, jags, and stan*. Waltham, MA: Academic Press.

Laming, D. R. J. (1968). *Information theory of choice-reaction times*. London: Academic Press.

Lee, M. D., & Wagenmakers, E.-J. (2013). *Bayesian cognitive modeling: A practical course*. Cambridge University Press.

Levav, J., & Fitzsimons, G. J. (2006). When questions change behavior: The role of ease of representation. *Psychological Science*, 17(3), 207–213.

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, 1171.

Luce, R. D. (1986). *Response times*. New York: Oxford University Press.

Massar, S. A., Lim, J., Sasmita, K., & Chee, M. W. (2016). Rewards boost sustained attention through higher effort: A value-based decision making approach. *Biological Psychology*, 120, 21–27.

Michael, J. R., Schucany, W. R., & Haas, R. W. (1976). Generating random variates using transformations with multiple roots. *The American Statistician*, 30(2), 88–90.

Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In *Proceedings of the 3rd international workshop on distributed*

statistical computing.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85, 59–108.

Ratcliff, R. (1979). Group reaction time distributions and an analysis of distribution statistics. *Psychological Bulletin*, 86, 446–461.

Ratcliff, R. (1993). Methods for dealing with reaction time outliers. *Psychological Bulletin*, 114, 510–532.

Ratcliff, R., Thapar, A., & McKoon, G. (2001). The effects of aging on reaction time in a signal detection task. *Psychology and Aging*, 16, 323–341.

Reddi, B., & Carpenter, R. (2003). Accuracy, information and response time in a saccadic decision task. *Journal of Neurophysiology*, 90, 3538–3546.

Rouder, J. N. (2005). Are unshifted distributional models appropriate for response time? *Psychometrika*, 70, 377–381.

Rouder, J. N., & Haaf, J. M. (in press). A psychometrics of individual differences in experimental tasks. *Psychonomic Bulletin and Review*.

Rouder, J. N., Pratte, M. S., & Morey, R. D. (2010). Latent mnemonic strengths are latent: A comment on Mickes, Wixted, and Wais (2007). *Psychonomic Bulletin and Review*, 17, 427–435.

Rouder, J. N., Province, J. M., Morey, R. D., Gomez, P., & Heathcote, A. (2015). The lognormal race: A cognitive-process model of choice and latency with desirable psychometric properties. *Psychometrika*, 80, 491–513.

Rouder, J. N., Tuerlinckx, F., Speckman, P. L., Lu, J., & Gomez, P. (2008). A hierarchical approach for fitting curves to response time measurements. *Psychonomic*

Bulletin & Review, 15(1201-1208).

Schwarz, W. (2001). The ex-Wald distribution as a descriptive model of response times. *Behavioral Research Methods, Instruments, and Computers*, 33, 457–469.

Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, 171, 701–703.

Simon, J. R. (1969). Reactions toward the source of stimulation. *Journal of Experimental Psychology*, 81, 174–176.

Sternberg, S. (1966). High-speed scanning in human memory. *Science*, 153, 652–654.

Stone, M. (1960). Models for choice-reaction time. *Psychometrika*, 25, 251–260.

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, 18, 643–662.

Thiele, J. E., Haaf, J. M., & Rouder, J. N. (2017). Bayesian analysis for systems factorial technology. *Journal of Mathematical Psychology*, 81, 40–54.

Ulrich, R., & Miller, J. O. (1993). Information processing models generating lognormally distributed reaction times. *Journal of Mathematical Psychology*, 37, 513–525.

Usher, M., & McClelland, J. L. (2001). On the time course of perceptual choice: The leaky competing accumulator model. *Psychological Review*, 108, 550–592.

Van Zandt, T. (2000). How to fit a response time distribution. *Psychonomic Bulletin and Review*, 7, 424–465.

Vickers, D., & Smith, P. (1985). Accumulator and random-walk models of

psychophysical discrimination: A counter-evaluation. *Perception*, 14, 471–497.

Wagenmakers, E. J., & Brown, S. (2007). On the linear relation between the mean and the standard deviation of a response time distribution. *Psychological Review*, 114, 830–841.

Whelan, R. (2008). Effective analysis of reaction time data. *The Psychological Record*, 58(3), 475–482.

Wilcox, R. R. (1998). How many discoveries have been lost by ignoring modern statistical methods? *American Psychologist*, 53(3), 300.

Figure captions

- Figure 1.* Untransformed and Transformed Response Times Distributions. There are four distributions in each panel forming two pairs of highly similar distributions. For each pair, one distribution denotes a baseline condition, the other a treatment condition. Distributions with solid lines are unshifted; those with dashed lines are highly shifted. **A.** Untransformed inverse Gaussian distributions of RT. Each pair has an effect in drift rate. **B** The distributions that result from the logarithm transforms. **C** The distributions that result from the reciprocal transform.
- Figure 2.* **A.** Estimated shifts from Gomez et al. (2008). **B.** An example of the accumulation in a one-bound diffusion model.
- Figure 3.* ROC plots for each of the canonical cases. Left column: fixed shift, right column: normally distributed shift.
- Figure 4.* The effect of adding a small degree of variability to the shift parameter is to make the left tail more gentle.

Table 1
Parameter Values For Simulations

Name	Avg (s)	Effect (s)	Shift	Bound	Drift ₁	Drift ₂	Drift ₀
Perception	0.30	0.02	0.15	1.20	7.50	8.57	8.04
Attention	0.65	0.05	0.35	1.20	3.69	4.36	4.03
Linguistics	1.35	0.15	0.35	1.20	1.12	1.30	1.21

Table 2
Power from Simulation I (fixed shifts)

			Power			Level		
	N	Shift	RT	log(RT)	1/RT	RT	Log(RT)	1/RT
Perception								
Realistic	60	0.15	0.61	0.61	0.59	0.050	0.049	0.050
Realistic	20	0.15	0.24	0.24	0.24	0.050	0.049	0.050
Realistic	120	0.15	0.89	0.89	0.88	0.050	0.049	0.050
Zero	60	0.00	0.61	0.59	0.54	0.050	0.049	0.050
Zero	20	0.00	0.24	0.24	0.22	0.049	0.050	0.049
Zero	120	0.00	0.89	0.88	0.83	0.050	0.050	0.051
Large	60	0.30	0.61	0.61	0.60	0.049	0.051	0.050
Large	20	0.30	0.24	0.24	0.24	0.049	0.049	0.050
Large	120	0.30	0.89	0.89	0.89	0.050	0.050	0.049
Attention								
Realistic	60	0.35	0.50	0.50	0.48	0.049	0.049	0.051
Realistic	20	0.35	0.19	0.20	0.19	0.049	0.049	0.050
Realistic	120	0.35	0.81	0.80	0.78	0.050	0.050	0.050
Zero	60	0.00	0.50	0.48	0.39	0.051	0.050	0.050
Zero	20	0.00	0.19	0.19	0.16	0.047	0.050	0.047
Zero	120	0.00	0.81	0.77	0.66	0.050	0.048	0.051
Large	60	0.70	0.50	0.50	0.49	0.050	0.051	0.050
Large	20	0.70	0.19	0.20	0.19	0.049	0.049	0.049
Large	120	0.70	0.80	0.80	0.79	0.049	0.051	0.051
Linguistic								
Realistic	240	0.35	0.50	0.46	0.36	0.049	0.050	0.050
Realistic	80	0.35	0.20	0.18	0.15	0.049	0.050	0.049
Realistic	480	0.35	0.80	0.74	0.62	0.049	0.050	0.050
Zero	240	0.00	0.50	0.42	0.25	0.049	0.050	0.050
Zero	80	0.00	0.20	0.17	0.12	0.049	0.048	0.048
Zero	480	0.00	0.80	0.70	0.44	0.049	0.050	0.050
Large	240	0.70	0.50	0.47	0.40	0.050	0.050	0.050
Large	80	0.70	0.20	0.19	0.17	0.049	0.050	0.050
Large	480	0.70	0.80	0.76	0.68	0.050	0.051	0.049

Table 3
Effect Sizes

Shift	RT	log(RT)	1/RT
Perception			
0.150	0.417	0.416	0.410
0.000	0.417	0.409	0.381
0.300	0.417	0.417	0.414
Attention			
0.350	0.368	0.368	0.360
0.000	0.368	0.357	0.314
0.700	0.368	0.369	0.366
Linguistic			
0.350	0.181	0.173	0.150
0.000	0.181	0.164	0.119
0.700	0.181	0.176	0.161

Table 4
Effect Sizes, Variable Shift

Shift	RT	log(RT)	1/RT
Perception			
0.150	0.380	0.374	0.360
0.300	0.311	0.303	0.291
Attention			
0.350	0.330	0.322	0.306
0.700	0.279	0.269	0.253
Linguistic			
0.350	0.177	0.169	0.147
0.700	0.176	0.170	0.154

Table 5
Power from Simulation II (random shifts)

			Power			Level		
	N	Shift	RT	log(RT)	1/RT	RT	log(RT)	1/RT
Perception								
Realistic	60	0.15	0.53	0.51	0.49	0.049	0.049	0.049
Realistic	20	0.15	0.21	0.20	0.19	0.049	0.051	0.050
Realistic	120	0.15	0.82	0.81	0.78	0.049	0.050	0.050
Large	60	0.30	0.37	0.36	0.34	0.051	0.049	0.049
Large	20	0.30	0.15	0.15	0.14	0.049	0.051	0.050
Large	120	0.30	0.64	0.62	0.59	0.049	0.050	0.049
Attention								
Realistic	60	0.35	0.45	0.43	0.40	0.050	0.050	0.050
Realistic	20	0.35	0.17	0.17	0.16	0.049	0.052	0.049
Realistic	120	0.35	0.75	0.72	0.68	0.050	0.050	0.049
Large	60	0.70	0.34	0.33	0.30	0.050	0.049	0.050
Large	20	0.70	0.14	0.13	0.13	0.049	0.050	0.050
Large	120	0.70	0.60	0.57	0.53	0.049	0.049	0.051
Linguistic								
Realistic	240	0.35	0.50	0.45	0.35	0.049	0.050	0.050
Realistic	80	0.35	0.20	0.18	0.15	0.050	0.050	0.050
Realistic	480	0.35	0.79	0.74	0.61	0.050	0.050	0.049
Large	240	0.70	0.50	0.45	0.37	0.050	0.050	0.050
Large	80	0.70	0.19	0.18	0.16	0.049	0.051	0.050
Large	480	0.70	0.79	0.74	0.64	0.049	0.049	0.051

Figure 1.

Figure 2.

Figure 3.

Figure 4.