Scale Invariant Contrasts of Response Latency Distributions

N. Sriram
University of Virginia

Brian A. Nosek
University of Virginia

Anthony G. Greenwald
University of Washington

Contact information:

N. Sriram
Box 400400; 102 Gilmer Hall
Department of Psychology
University of Virginia
Charlottesville, VA 22904
E-mail: nsriram@virginia.edu
Word count = 11739

Abstract

Individual differences in general speed lead to a positive correlation between the mean and standard deviation of mean latency. This “coarse” scaling effect causes the mean latency difference (MLD) to be spuriously correlated with general speed. Within individuals, the correlation between the mean and standard deviation of trial latencies leads contrasted distributions to increase their overlap as an MLD of fixed width is translated to the right. To address this “fine” scaling effect, contrasts based on within subject latency transformations including the logarithm, standardization, and ranking were evaluated and turned out to be distinctly superior to the MLD. Notably, the mean gaussian rank latency difference was internally consistent, eliminated fine scaling, meliorated coarse scaling, reduced correlations with general speed, increased statistical power to detect within subject and between group effects, and has the potential to increase the validity of inferences drawn from response latency data.

Abstract = 145 words

Keywords: Latency Difference, Scale Invariance, Ranking, Standardization, Gaussian Latency

The mean response latency difference, $\overline{L}_A$, is a measure of the processing demands in one task, relative to a contrasted task (Meyer, Osman, Irwin, & Yantis, 1988). This comparison operationalizes effects for experimental paradigms across psychology including semantic priming (Neely, 1991), lexical decisions (Balota & Chumbley, 1984), evaluative priming (Fazio, Sanbonmatsu, Powell, & Kardes, 1986), the Implicit Association Test (Greenwald, McGhee, & Schwartz, 1998), the Stroop task (MacLeod, 1991), and encompassing the vast majority of speeded binary classification tasks (Proctor & Cho, 2006). In each, multiple response latencies for individual trials are averaged then a difference between means for two conditions ($\overline{L}_A$) creates the “coarse” unit of analysis – the summary score for comparative task performance.

Despite widespread use, $\overline{L}_A$ has systematic psychometric biases that seriously compromise inferences drawn from it. These biases emerge from associations between mean and variance at two distinct levels of analysis. At the coarse level, where mean latency is the unit of analysis, there is a positive correlation between the mean and standard deviation of mean latencies across tasks. In a sample, tasks generally associated with longer mean latencies also have a larger variance of mean latency across subjects. We label this correlation as the *coarse scaling* effect.

The “fine” unit of analysis is the raw latencies of the individual response trials within a task. At the fine level, there is a positive correlation between the mean and standard deviation of raw response latencies. Within subjects across tasks and within tasks across subjects, the mean and standard deviation of latency distributions are positively correlated. We label this correlation as the *fine scaling* effect and it is conceptually distinct from the coarse scaling effect.

Both coarse scaling and fine scaling reflect correlations between mean and standard deviation. However the former is at a “coarser” level of granularity (mean latency) whereas the latter is at a “finer” level of granularity (raw latency). Various treatments of response latency distributions comparisons have considered one or the other, but, to our knowledge, never both simultaneously.

The psychometric biases in the mean response latency difference emerge from coarse scaling and fine scaling effects. Sriram, Greenwald, and Nosek (2006) delineate the effects of coarse scaling on $\overline{L}_A$. Coarse scaling causes $\overline{L}_A$ to be asymmetrically influenced by the larger and more variable mean latency. Across subjects, $\overline{L}_A$ is constrained to be negatively correlated with general processing speed, an index of cognitive fluency. That is, subjects who are generally faster have smaller $\overline{L}_A$ than subjects who are generally slower. Coarse scaling also attenuates correlations of $\overline{L}_A$ with convergent criteria. To mitigate the coarse scaling effect, Sriram et al. (2006) evaluated transformations of mean latency but, while they varied in their effectiveness, they failed to meliorate the biases adequately.

In the rate amount framework (Faust, Balota, Spieler, & Ferraro, 1999), tasks vary in their information processing amount, individuals differ in their information processing rate, and mean response latency is the product of information amount and the reciprocal of information processing rate. The rate amount framework provides a general, task independent, account of variability in mean response latencies. Faust et al. (1999) show that the assumptions of the rate amount framework constitute sufficient conditions for coarse scaling and a simulation to this end is provided in this report.

Fine scaling complicates the interpretation of within-subjects $\overline{L}_\Delta$ comparisons because, for a fixed $\overline{L}_\Delta$, the overlap between response latency distributions increases with location on the scale; that is, as $\overline{L}_\Delta$ is translated to the right. For example, $\overline{L}_\Delta = 100$ ms has different implications depending on whether it is centered at 500 ms or 800 ms. Scale invariance is violated because the interpretation of $\overline{L}_\Delta$ depends on location on the scale. Fine scaling is a statistical property of response latency distributions. Fine scaling is shown to emerge from variation in information rate parameters embedded within cognitive models of response latency generation such as diffusion (Ratcliff, Gomez, & McKoon, 2004) and LATER (Carpenter, 1999).

To address the fine scaling effect, this report examines a variety of contrasts derived from transforming within-subject trial latencies. This evaluation includes the mean log latency difference $\log \overline{L}_\Delta$, mean rank differences $\overline{F}_\Delta$ and $\overline{G}_\Delta$, and mean standardized differences $\overline{C}_\Delta$, $\overline{d}_\Delta$, and $\overline{D}_\Delta$. In homogenizing variance of latencies across conditions within subjects, these transformations directly reduce fine scaling. In addition, they meliorate coarse scaling biases by indirectly homogenizing variance of condition contrasts across subjects. Simulations and empirical data analyses show that, compared to the mean latency difference, these contrasts exhibit greater statistical power to detect valid effects and show reduced confounding with general processing speed. The mean rank and mean standardized differences are particularly effective in this regard and are recommended for improving the validity of inferences drawn from response latency data.

Sources of Coarse and Fine Scaling

Coarse Scaling in the Rate Amount Framework

In the rate amount framework tasks vary in their information processing amount and individuals differ in their information processing rate or cognitive speed (Faust et al., 1999). The mean latency for a subject is $\kappa + \frac{\delta}{\tau}$ where κ is constant within individual, δ is task information processing amount, and τ is individual cognitive speed. A simulation, involving 1000 samples with 100 subjects in conditions x and y , examined the implications of the rate amount model for

$\lambda = \frac{\sigma_y}{\sigma_x}$, the ratio of the standard deviation of mean latencies and r_{xy} , the correlation between

mean latencies. In particular, the simulation examined the dependence of λ and r_{xy} on three independent variables that were varied across samples: (i) the effect size contrast, (ii) measurement error in mean latency, and (iii) variability in cognitive speed across conditions.

The mean latencies in two conditions, x and y , were obtained from parameters that were assigned as follows. For each subject, κ and τ were sampled from a normal distribution, $\mu = 100$ ms $\sigma = 20$ ms, truncated at 20 ms and $\mu = .0025$ bits/ms $\sigma = .0005$ bits/ms, truncated at .0005 bits/ms. Modest random within subject variability in τ across conditions was introduced so that for each sample, σ_τ was selected from [1, 10]% of τ . For each subject, τ_x and τ_y were selected from a normal distribution, $\mu = \tau$ and $\sigma = \sigma_\tau$. The mean correlation between τ_x and τ_y , across samples, was less than 1 ($r = .92$) because τ varied a little across conditions.

To vary the size of the condition contrast across samples, δ_x was constant at 1 bit while δ_y was randomly selected from [1, 2]. The standard deviation of gaussian measurement error in mean latencies, σ_e was varied across samples to be [20, 100] ms. At an average of 60 ms, this measurement error is substantial, given that mean latencies are aggregated indices. The condition mean latencies for subjects were assigned using $x = \kappa + \frac{\delta_x}{\tau_x} + e_x$ and $y = \kappa + \frac{\delta_y}{\tau_y} + e_y$.

The grand sample means for x and y were 520 ms (SD=12) and 741 ms (SD=120) respectively. Coarse scaling was reflected in the .76 correlation between $\bar{y}$ and σ_y across samples and implies, with high probability, that $\lambda > 1$. For these data $\lambda > 1$ in 955 of 1000 samples, $\bar{\lambda} = 1.38$, SD=0.25. The mean δ_y was 1.49 (SD=.28) and the mean r_{xy} , at .735 (SD=0.13), comparable to empirical values (Faust et al., 1999, p.781). The assumptions in the rate amount model constitute sufficient conditions for coarse scaling as they lead to a positive correlation between the mean and standard deviation of mean latency.

How do λ and r_{xy} vary as a function of contrast magnitude δ_y , measurement error σ_e , and variability in cognitive speed σ_τ ? These independent variables were uncorrelated with each other. λ was strongly associated with δ_y ($r = .87$), had a weaker negative relation to σ_e ($r = -.215$), and uncorrelated with σ_τ ($r = .05$). Thus λ , identified as the principal cause of correlational biases in $\bar{L}_\Delta$ (Sriram et al., 2006), strongly depends on effect size and is reduced by measurement error. r_{xy} is unrelated to effect size δ_y ($r = .08$) and is negatively related to σ_e ($r = -.843$) and σ_τ ($r = -.216$). If measurement error is reduced, then r_{xy} has a strong negative association with σ_τ . From the rate amount perspective, correlations between mean latencies cannot be ascribed to differences in information amounts across tasks. Rather, they depend on random measurement error in mean latency and random intra-individual variability in cognitive speed. If measurement error is minimal, then typical values of r_{xy} (e.g., .7) must be taken to indicate some variation in individual cognitive speed across tasks if the values of information amount are taken to be task associated constants.

Coarse Scaling in Brinley Plots

Figure 1 is a Brinley scatter plot that depicts the covariation of sample average mean latencies, across a variety of tasks, between younger and older samples (Spieler, 2001). The strong linear relationship, accounting for over 90% of the scatter plot variance, implies that task performance in younger and older samples have the same rank order. However, task mean latencies are both larger and more variable in older than younger samples. Given the nearly perfect linear fit, the regression slope is numerically equal to the ratio of the standard deviations of the task mean latencies across the older and younger samples (Myerson, Adams, Hale, & Jenkins, 2003). The slope ϕ is termed as the age related slowing factor, is usually around 1.5, but can vary across disparate task domains (Verhaeghen et al., 2002).

The pattern in the two-sample Brinley plot can be generalized to the performance of multiple samples on a variety of tasks. If samples vary in their average general speed, then the mean and variance of the sample mean latency across tasks would be correlated very highly. Samples that have lower average general speed would have larger means and larger variances

across tasks. This positive mean-variance correlation is similar to within-sample coarse scaling but emerges at an even higher, coarser, level of aggregation.

Coarse Scaling Effects on $\overline{L}_A$ Comparisons

Sriram et al. (2006) showed that the greater the departure of λ from 1, the higher the correlation between $\overline{L}_A$ and the average of the two mean latencies. In the presence of coarse scaling, subjects with greater average mean latencies have larger values of $\overline{L}_A$. This makes inferences from $\overline{L}_A$ comparisons across subjects problematic because a larger $\overline{L}_A$ may largely reflect lower general speed. From the perspective of the rate amount model, the difference between the $\overline{L}_A$ of two individuals A and B is $(\delta_Y - \delta_X) \left[\frac{1}{\tau_A} - \frac{1}{\tau_B} \right]$. For a given task contrast, $\delta_Y - \delta_X$ is constant and the $\overline{L}_A$ difference varies directly as the difference between the reciprocals of general speeds. The ratio between the $\overline{L}_A$ of A and $\overline{L}_A$ of B is $\frac{\tau_B}{\tau_A}$ can be viewed as a between-individuals scaling factor that leads to spurious effects.

In the Stroop task, $\overline{L}_A = 150$ ms is interpreted as a larger interference effect than $\overline{L}_A = 100$ ms. This inference assumes that both mean differences are located on a common interval scale (Stevens, 1946). The literature on cognitive aging questions this scaling assumption in between-groups comparisons of $\overline{L}_A$ (Salthouse, 1985; Salthouse & Hedden, 2002). In Brinley plots with almost perfect linear fits, the ratio of task contrasts across groups A and B, $\frac{\overline{L}_{AB}}{\overline{L}_{AA}} = \phi$ the slope of the regression line. This ratio is the same across *any* pair of tasks, and the magnitude of $\overline{L}_{AB} - \overline{L}_{AA}$ is directly proportional to ϕ . Like the between-individuals scaling ratio $\frac{\tau_B}{\tau_A}$, the between-groups scaling factor ϕ creates spurious differences in $\overline{L}_A$ across groups (Verhaeghen & Cerella, 2002).

Fine Scaling Effects in Individual Latency Distributions

The ex-gaussian distribution has been applied extensively to model response latency distributions and is specified using gaussian parameters μ , σ and an exponential parameter τ . The mean of the ex-gaussian is $\mu + \tau$ and the standard deviation is $\sqrt{\sigma^2 + \tau^2}$ and the common exponential parameter constrains them to be positively correlated. Fine scaling is present in alternative statistical models of latencies including the lognormal, displaced gamma, and the reciprocal gaussian distributions (Luce, 1986; Heathcote, Brown, & Mewhort, 2002). Fine scaling is a consequence of parameters shared between the mean and standard deviation. In the lognormal distribution, ζ and σ are location and scale parameters in the mean $= e^{\zeta + \frac{\sigma^2}{2}}$ and standard deviation $= \sqrt{e^{2\zeta + \sigma^2} (e^{\sigma^2} - 1)}$ (Patel, Kapadia, & Owen, 1976).

Fine scaling is also implied by process models such as diffusion (Ratcliff & Rouder, 1998) and LATER (Carpenter, 1999). The diffusion model is a continuous random walk sequential sampling model for binary classification tasks. In this model there is a random

accumulation of information that eventually reaches either one of two response thresholds. The drift rate is related to the mean speed with which information is accumulated and usually ranges from .1 to .5. The diffusion function has a variance that acts like a scaling parameter and is set to .01. The random walk commences from a starting point that is midway between the decision thresholds for two equally plausible alternative responses whose separation ranges from .07 to .17. Values of low drift rates are associated with higher mean latencies and errors as the diffusion process could occasionally reach the “wrong” threshold. Wagenmakers, Grasman, and Molenaar (2005) used these basic parameters and ignored more complex aspects used to fit empirical distributions like inter-trial variability in drift rates and starting point. For a plausible range of drift rates in the diffusion model, the mean and standard deviation of latency distributions were linearly related. Although this analysis did not vary drift rate across trials, the diffusion process suffices to account for the variability in response times. That is, the within-trial standard deviation of the drift rate in a diffusion model is a sufficient condition for fine scaling.

In the Linear Approach to Threshold with Ergodic Rate model (LATER; Carpenter, 1999), information systematically accumulates in a linear fashion within a trial to reach the response threshold. However, and crucially, the rate of information rise varies randomly across trials. According to Carpenter, the randomness of reaction time is due to a deliberate, gratuitous neural process that confers biological advantages to the organism executing the response. The reciprocal of latency follows a normal distribution and the latencies follow the reciprocal gaussian distribution. LATER conceptualizes the rate of information rise to be distributed normally with a particular mean and inter-trial standard deviation. In a given trial, the response latency is proportional to the reciprocal of the rate of information rise.

LATER was simulated on 1000 subjects with 100 trials each. Each subject was assigned a mean drift rate from [.002, .001]. The standard deviation of rate of information rise, across trials, was a constant .0002 for all subjects. The rate on a particular trial was selected from the corresponding normal distribution. Gaussian measurement error with SD=100 ms was added to the reciprocal of the rate to obtain the trial latency. The overall mean RT was 712 ms and the mean SD was 153 ms. The correlation between the mean and SD of RT was .92 and clearly indicates fine scaling.

Fine Scaling in Within-Subject $\bar{L}_A$ Comparisons

Consider single subject data from a 2 x 2 experiment in which simple and choice reaction tasks were administered under conditions of low and high alertness. For these simulated data, the following ex-gaussian parameters were used to generate latency distributions:

$\tau = .25 \mu$, $\sigma = .125 \mu$. The distributions for the Simple_{high} (M=500, SD=112), Choice_{high} (M=600, SD=134), Simple_{low} (M=900, SD=201), and Choice_{low} (M=1000, SD=224) conditions are displayed in Figure 2. In each condition, 80% of the variance is due to the positively skewed exponential component. The mean and SD have the same rank order across the conditions, a consequence of the linear relationship between the mean and standard deviation in ex-gaussians.

Fine scaling makes the direct interpretation of within-subject differences in $\bar{L}_A$ problematic. The mean-variance association governs the relationship between the location of $\bar{L}_A$ and the overlap between task latency distributions. In Figure 2, the location of the low alertness $\bar{L}_A = 100$ ms is at 900 ms whereas the high alertness $\bar{L}_A = 100$ ms is at 550 ms. Although $\bar{L}_A$ is constant across alertness levels, the overlap between distributions is far greater in the low

alertness condition (Cohen's $d = 0.47$) than in the high alertness condition ($d = 0.81$). The dominance probability that latencies in the Choice task are greater than latencies in the Simple task is 0.744 in the high alertness condition, much greater than the 0.648 in the low alertness condition. For completely overlapping distributions, $d=0$ and the dominance probability is 0.5.

In summary, while the distributions in the low and high alertness conditions have equidistant means, they differ strongly in the degree to which the distributions overlap. In the faster condition, there is bigger separation of the simple vs. choice distributions than in the slower condition, as illustrated by Cohen's d and dominance probabilities. The zero interaction conclusion using mean latencies is a consequence of fine scaling. Contrasts using Cohen's d and dominance probability reveal that the alertness manipulation has an increased effect at high levels of alertness. A proper consideration of the marginal difficulty by alertness in choice versus simple reaction time should incorporate information about the increased variability in responses when response latencies are slower; this information is not incorporated in $\overline{L}_A$.

Violation of Scale Invariance

To illustrate the general nature of the dependence of distribution separation on scale location a $5 \times 5 \times 5 \times 5$ simulation that varied the form of the ex-gaussians across a total of 6250 simulated subjects was conducted. The factors were $\overline{L}_A$ (20, 40, 60, 80 or 100 ms), the average of the two mean latencies or *position* (500, 625, 750, 875 or 1000 ms), *percent exponential* τ in the mixture (10, 20, 30, 40 or 50 %), and σ as a percentage of μ , or *percent gaussian* σ (5, 10, 15, 20, or 25 %). For every subject, 100 latencies were sampled from each of the two contrasted ex-gaussian distributions and latencies less than 100 ms were truncated.

Across 6250 subjects, the two measures of distribution separation, Cohen's d and dominance probability, were highly correlated at .966. Their correlations with the sample $\overline{L}_A$ were much more modest at .684 and .673 respectively. That is, over 50% of the variance in distributional measures of separation was distinct from $\overline{L}_A$ variance. A multiple regression including $\overline{L}_A$ magnitude and scale location as independent variables predicted 54% of the variance in Cohen's d , with a positive coefficient for $\overline{L}_A$ magnitude ($b = 4.02 \times 10^{-3}$, $\beta = .686$) and a negative coefficient for scale location ($b = -4.28 \times 10^{-4}$, $\beta = -.274$). As scale location increased, contrasted distributions with constant $\overline{L}_A$ had increased overlap. Within a subject, the effect of scale location on degree of overlap would be much stronger than this simulation indicates because within-subject ex-gaussians across conditions would arise from similar distributional parameters.

Scaling and Parameter Estimation

The antecedents of coarse scaling and fine scaling were identified in terms of theoretical parameters in models of mean latency and response latency distributions. Coarse scaling can be viewed as an outcome of individual differences in cognitive speed. Fine scaling emerges from intra-trial variability in drift rate in the diffusion model and inter-trial variability in linear rate of rise in the LATER model.

The estimation of cognitive speed and drift rate parameters from latency data with limited number of trials per condition is complex, sensitive to model specification, and constitutes a challenging problem. To ease parameter estimation and enhance interpretability, Faust et al.'s (1999) rate amount model of mean latencies excludes interactions between conditions and

subjects. That is, the possibility that individual cognitive speed might vary markedly across tasks was excluded from consideration.

Interactions between conditions and subjects are, however, at the heart of dissociations observed, for example, in social cognition research where liberals and conservatives could display negatively sloped Brinley plots; tasks tapping political attitudes may include some that are harder for conservatives but easier for liberals and others that are easier for conservatives but harder for liberals, such as one task that measures association strengths of *George Bush* with *good* and *John Kerry* with *bad* versus another that measures the association strengths of *John Kerry* with *good* and *George Bush* with *bad* (Greenwald, Nosek, & Banaji, 2003; Nosek, Banaji, & Greenwald, 2002; Nosek, et al., 2006). Faust et al. (1999) support their preference for simplicity by pointing out that well defined, complex models (e.g., Fisher & Glaser, 1996) are accompanied by strong theoretical assumptions that may not be readily applicable to diverse tasks.

Despite the simplicity of their model, Faust et al. (1999) sound a note of caution regarding the sensitivity of their model parameter estimates to the range of mean latencies across conditions. Models that jointly estimate person and condition parameters necessarily depend on the vagaries of the particular sample; changing the sample affects person estimates even though individual data are unchanged. For these reasons, Faust et al. (1999) recommend within subject data transformations that do not use estimation. This article follows this approach and examines a variety of scale invariant within subject latency transformations.

Scale Invariant Contrasts

Three classes of transformations are examined in this article. The logarithm transformation, recommended to homogenize variance when the mean and standard deviation are correlated (Snedecor & Cochran, 1989), is the basis for $\overline{\log L_A}$. The second type of transformation is based on rescaling $\overline{L_A}$ by dividing it with a specific type of within-subject standard deviations. The standardized measures include $\overline{d_A}$ (Cohen, 1988), $\overline{D_A}$ (Bush, Hess, & Wolford, 1993; Greenwald, Nosek, & Banaji, 2003), and $\overline{c_A}$ (Faust et al., 1999). The third class of measures are rank statistics, the mean fractional rank difference $\overline{F_A}$ and the mean gaussian rank difference $\overline{G_A}$. $\overline{F_A}$ is linearly related to the probability with which latencies from one within subject latency distribution dominate latencies from a contrasting distribution and $\overline{G_A}$ is a variant that is based on transforming fractional ranks to a gaussian distribution. $\overline{\log L_A}$, $\overline{d_A}$, $\overline{D_A}$, and $\overline{c_A}$ are not associated with particular measurement units are invariant to linear transformations of latencies. However, they are not invariant to non-linear monotone transformations. $\overline{F_A}$ and $\overline{G_A}$ are unit free and invariant to both linear and non-linear monotone data transformations. The term scale invariant is used to qualify all these transformations and refers to unit-free measures that are invariant under linear transformations.

The Logarithm Transformation

At the level of trial latencies, the median (Miller, 1988) and reciprocal latency transformations (Greenwald, Nosek, & Banaji, 2003) reduce the validity of latency data analysis. The log transformation of individual latencies is prescribed to alleviate skewed latencies and the confounding of the latency difference with general processing speed (Meiran, 1996) and is particularly appropriate when the mean and standard deviation are correlated (Snedecor &

Cochran, 1989, pp. 286-287). The mean log latency difference $\overline{\log L_{\Delta}}$ is the ratio of the logarithms of the geometric means of latencies in two contrasted conditions. This measure can reduce the role of individual differences in processing speed across conditions. If latencies in each condition were scaled versions of general processing speed, the ratio would eliminate the confounding caused by processing speed. Faust et al. (1999, p.794) note however that empirical latency data do not, in general, satisfy this assumption. Bush et al. (1993) found that in their simulation studies that, in terms of effect size sensitivity, the logarithm was noticeably inferior to standardizing trial latencies, particularly when the latency distributions were symmetric.

Standard Deviation Transformations

For repeated measures ANOVA of mean latencies, Faust et al. (1999, pp. 788-789) recommend a standardizing transformation, motivated by their rate amount model, that normalizes each participant's condition mean latency by dividing it by $\sigma_{between}$, the standard deviation of the condition mean latencies. If this transformation is applied to the case of two

conditions $\sigma_{between} = \frac{\overline{L_{\Delta}}}{2}$ and the standardized latency difference has zero variance across subjects. This standardization may be applied when data is partitioned into three or more conditions and the transformed latency is labeled c (condition) and $\overline{c_{\Delta}} = \frac{\overline{L_{\Delta}}}{\sigma_{between}}$.

Cohen's d (1988) uses the average within condition standard deviation as the normalizing factor and can be extended to the case of two or more conditions and the contrast is labeled as

$\overline{d_{\Delta}} = \frac{\overline{L_{\Delta}}}{\sigma_{within}}$. σ_{within} and $\sigma_{between}$ jointly determine the extent to which the condition latency

distributions overlap each other. The overlap is minimal if σ_{within} is small and $\sigma_{between}$ is large and the overlap is extensive when σ_{within} is large and $\sigma_{between}$ is small.

In the particular case of two conditions, Faust et al. (1999, p.791) suggest using the standard deviation of condition trial latencies pooled together. This standardization was applied in simulations by Bush et al. (1993) and found to be superior to the logarithm and other alternatives. This strategy was evaluated empirically on large data sets by Greenwald, Nosek, and Banaji (2003) and forms the basis for the within-subjects D score contrasting two conditions. In addition to standardization, the D score algorithm eliminates extreme outliers, treats error responses, and eliminates participants with very many fast, anticipatory responses. For the data that Greenwald et al. evaluated, 10000 ms was set as the threshold for eliminating observations and trials had a built-in error penalty as subjects were obliged to provide the correct response. For tasks without built-in error penalties, they recommend replacing error latencies with a value that is 1 SD higher than the condition mean. To keep the terminology consistent, the standard

deviation of all latencies, σ_{total} is the scaling factor such $\overline{D_{\Delta}} = \frac{\overline{L_{\Delta}}}{\sigma_{total}}$.

In the case of two conditions, Cohen's d and the D score are identical to $\overline{d_{\Delta}}$ and $\overline{D_{\Delta}}$ respectively. $\overline{D_{\Delta}}$ is identical to the standardization applied by Bush et al. (1993) who termed it the Z transformation. This article uses D to emphasize the relatedness between the three types of standardization. If latency data is partitioned into k within subjects conditions, c , d , and D refer

to latencies divided by the appropriate standard deviation. Given that $\sigma_{total}^2 = \sigma_{between}^2 + \sigma_{within}^2$, where the variances are the respective sums of squares divided by the total number of latencies, $\overline{d_\Delta}$ and $\overline{c_\Delta}$ are expected to be relatively distinct. However $\overline{D_\Delta}$ would have greater similarity with $\overline{d_\Delta}$ and $\overline{c_\Delta}$ as it considers both variances in its standardization. These measures share the sign of $\overline{L_\Delta}$, and each one of them can be expressed in terms of the other two.

$$\overline{c_\Delta} = \frac{\overline{d_\Delta} \overline{D_\Delta}}{\sqrt{\overline{d_\Delta}^2 - \overline{D_\Delta}^2}}, \quad \overline{d_\Delta} = \frac{\overline{c_\Delta} \overline{D_\Delta}}{\sqrt{\overline{c_\Delta}^2 - \overline{D_\Delta}^2}}, \quad \overline{D_\Delta} = \frac{\overline{c_\Delta} \overline{d_\Delta}}{\sqrt{\overline{c_\Delta}^2 + \overline{d_\Delta}^2}} \quad (1)$$

Rank Transformations

What are the consequences of ranking within subject latencies? The rank transformation is invariant to monotone transformations, and consequently, is unaffected by any change of scale that preserves the rank order of the within-subject trial latencies such as the square root, reciprocal, or logarithm (Cliff, 1996). Ranking transforms within subject latencies to a uniform distribution. The influence of outliers is reduced as every observation contributes equally to the mean rank.

Fractional rank latency difference. Let F_i be the fractional rank of L_i in the ranking of latencies across k conditions. F_i ranges from a minimum of $\frac{1}{\sum_{i=1}^k n_i}$ to a maximum of 1, where

$N = \sum_{i=1}^k n_i$ is the total number of trial latencies. For $k=2$, the mean fractional rank difference $\overline{F_\Delta}$ is a scale invariant index that contrasts latencies in two conditions, x and y . The ranks of observations in each condition are averaged and the difference is divided by $n_x + n_y$ to represent it as a proportion. $\overline{F_\Delta}$ ranges from -0.5 to 0.5, with zero representing no difference and ± 0.5 indicating no overlap between the two response distributions. $\overline{F_\Delta}$ is expressed in terms of the raw ranks R as follows:

$$\overline{F_\Delta} = \frac{1}{(n_x + n_y)} \left(\frac{\sum_{j=1}^{n_y} R_j}{n_y} - \frac{\sum_{i=1}^{n_x} R_i}{n_x} \right) \quad (2)$$

$$\sum_{m=1}^{n_y} m = \frac{n_y(n_y + 1)}{2} \quad \text{and} \quad \sum_{m=1}^{n_x + n_y} m = \frac{(n_x + n_y)(n_x + n_y + 1)}{2} \quad \text{are arithmetic identities.} \quad (3)$$

$$\text{Observe that } \sum_{i=1}^{n_x} R_i = \sum_{m=1}^{n_x + n_y} m - \sum_{j=1}^{n_y} R_j \text{ as the sum of the } n_x + n_y \text{ ranks is a constant.} \quad (4)$$

$$\text{Substituting (4) in (2), } \overline{F_\Delta} = \frac{2 \sum_{j=1}^{n_y} R_j - n_y(n_x + n_y + 1)}{2n_x n_y} \quad (5)$$

$\overline{F}_\Delta + 0.5$ is the mean probability with which observations in one distribution dominate those in a contrasting distribution. To prove the preceding assertion, we use the definition provided by Handcock and Morris (1999) in their monograph on relative distributions. Let latencies in conditions x and y have ranks, R , ranging from 1 to $n_x + n_y$. Latencies in condition y are ranked from 1 to n_y . The relative distribution index of the i^{th} latency in condition y , L_i , is termed $R_{\Delta i}$ and is the difference between its combined rank $R_{L_i}^{x \cup y}$ and its rank within y , $R_{L_i}^y$ divided by n_x .

$$R_{\Delta i} = \frac{R_{L_i}^{x \cup y} - R_{L_i}^y}{n_x} \quad (6)$$

$R_{\Delta i} = 0.75$ for a particular latency implies that it exceeds 75% of the values in condition x . Observations in the non-overlapping parts of contrasting distributions have $R_{\Delta i} = 1$ of either 1 or 0. The mean $R_{\Delta i}$ across n_y observations, $\overline{R}_\Delta$, represents the probability with which observations in y dominate those in x .

$$\text{Using (6) and (3), } \overline{R}_\Delta = \frac{1}{n_y} \left(\frac{\sum_{j=1}^{n_y} R_j - \sum_{m=1}^{n_y} m}{n_x} \right) = \frac{2 \sum_{j=1}^{n_y} R_j - n_y(n_y + 1)}{2n_x n_y} \quad (7)$$

$$\text{Comparing (5) and (7), } \overline{R}_\Delta = \overline{F}_\Delta + \frac{1}{2} \quad (8)$$

In the two group context, $\overline{F}_\Delta = 0$ implies that the mean rank in the two groups are identical or $\overline{F}_y = \overline{F}_x$. Values of .5 and -.5 indicate that latencies in y completely dominate those in x and vice versa respectively. $\overline{F}_\Delta$ is intrinsically a relative measure. The sum

$$\overline{F}_y + \overline{F}_x = \frac{2 \sum_{j=1}^{n_y} R_j (n_x - n_y)}{n_x n_y (n_x + n_y)} + \frac{n_x + n_y + 1}{n_y} \quad (9)$$

$$\text{If } n_x = n_y = n, \overline{F}_y + \overline{F}_x = 2 + \frac{1}{n} \quad (10)$$

The $\overline{F}_\Delta$ components, $\overline{F}_y$ and $\overline{F}_x$, are redundant. For constant n , they correlate -1 with each other and have identical variance across subjects. As the mean rank sum is constant, it is uncorrelated with the mean rank difference. $\overline{F}_\Delta$ measures distribution overlap and does not contains information that is unique to either distribution.

Gaussian rank latency difference. Bartlett (1947, p. 50), commenting on Friedman's (1937) analysis of ranks, remarked that, once data were ranked, expected normal scores could be used, in lieu of ranks. Such a transformation remains monotone invariant and, in his view, the potential to improve the validity of ANOVA conducted on data that fully meets the normality assumptions. This section details the computation of expected normal scores or the gaussian rank latency difference. In ranking, the distribution of within subject latencies is flattened into a

uniform distribution. Each latency is transformed into a fractional rank or percentile value relative to the latencies across all conditions. This percentile represents the probability with which a given value equals or exceeds other values in the distribution. The normal deviate with this cumulative probability represents the scale value of the observation in a standard gaussian latency distribution with $\mu = 500$, $\sigma = 100$ (these values are arbitrary and chosen for ease of interpretation). Like uniform fractional ranks, gaussian rank latencies are scale invariant. The percentiles 0.16, 0.5, and 0.84 correspond to gaussian rank latencies 400 ms, 500 ms, and 600 ms respectively.

In finite data, fractional ranks of zero are absent; the minimum value is assigned a rank which is the reciprocal of the sample size and the value of 1 is assigned to the maximum value.

To ensure that all percentiles have a finite normal deviate, $\frac{1}{2\sum_{i=1}^k n_i}$ or half the width between

successive F_i values is subtracted from each fractional rank prior to computing the normal deviate corresponding to the percentile. The distribution of the gaussian rank latencies is symmetric, centered around the mean of 500, and has a standard deviation of 100. For 100 observations, gaussian rank latencies can range from 242 ms to 758 ms, with the 50th and 51st ranked values being 498.75 ms and 501.25 ms respectively. The mean gaussian rank difference is $\overline{G}_\Delta$.

Cohen's d, D, and P(d)

In the case of two conditions, $\overline{D}_\Delta$ or the D score is the mean difference divided by the standard deviation of the trial latencies σ_{xy} . Each latency is subtracted from the grand

mean, $\frac{n_x \bar{x} + n_y \bar{y}}{n_x + n_y}$, and then divided by σ_{xy} . The total sums of squares of the trial latencies in the two conditions can be partitioned as follows.

$$SS_{total} = (n_x + n_y) \sigma_{xy}^2 = n_x \left(\bar{x} - \frac{n_x \bar{x} + n_y \bar{y}}{n_x + n_y} \right)^2 + n_y \left(\bar{y} - \frac{n_x \bar{x} + n_y \bar{y}}{n_x + n_y} \right)^2 + n_x \sigma_x^2 + n_y \sigma_y^2 \quad (11)$$

$$\sigma_{xy}^2 = \frac{n_x n_y (\bar{y} - \bar{x})^2}{(n_x + n_y)^2} + \frac{n_x \sigma_x^2 + n_y \sigma_y^2}{(n_x + n_y)} \quad (12)$$

Dividing throughout by $(\bar{y} - \bar{x})^2$ we get

$$\frac{\sigma_{xy}^2}{(\bar{y} - \bar{x})^2} = \frac{1}{D^2} = \frac{n_x n_y}{(n_x + n_y)^2} + \frac{n_x \sigma_x^2 + n_y \sigma_y^2}{(n_x + n_y)(\bar{y} - \bar{x})^2} \quad (13)$$

$$|\max D| = \frac{n_x + n_y}{\sqrt{n_x n_y}} \text{ when the last term in (13) vanishes} \quad (14)$$

The magnitude of the D score is positively associated with $\overline{L}_\Delta^2$ and inversely related to the sum of the within-group variances. These variances act in contrary ways in influencing the separation between contrasted distributions. The D score and Cohen's d are monotonic with respect to each other. Cohen's d has no upper bound whereas, for equal size groups, D assumes a maximum value of 2 as seen in (15). In the equal groups case, the D score is exactly twice the

point biserial correlation between a binary variable coding the conditions and the trial response latencies.

$$\frac{\sigma_{xy}^2}{(\bar{y} - \bar{x})^2} = \frac{1}{D^2} = \frac{1}{4} + \frac{1}{d^2}, \quad D = \frac{2d}{\sqrt{4 + d^2}}, \quad \text{Cohen's } d = \frac{2D}{\sqrt{4 - D^2}} \quad (15)$$

$$\text{Point Biserial Correlation} = \frac{(\bar{Y} - \bar{X})}{\sigma_{XY}} \frac{\sqrt{n_1 n_2}}{n_1 + n_2} = D \frac{\sqrt{n_1 n_2}}{n_1 + n_2} \quad (16)$$

For overlapping normal distributions with identical variance, the dominance probability $P(d) = \bar{R}_\Delta = \bar{F}_\Delta + \frac{1}{2}$ is a monotonic function of Cohen's d .

$$P(d) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \int_{-\infty}^y e^{-\frac{x^2}{2}} e^{-\frac{(y-d)^2}{2}} dx dy \quad (17)$$

This iterated integral can be understood by visualizing two overlapping normal distributions, left and right, with unit variance. The one on the left has a mean of zero and the other, to the right, has a mean of d . The inner integral, varying in x , is the cumulative probability for the left distribution at a given threshold y in $[-\infty, +\infty]$. The cumulative probability is the percentage of values in the left distribution that are less than or dominated by this threshold. As we desire the domination of the left by the right, the cumulative probability is weighted by the density of the right distribution. This weighting is implemented by the outer integral, varying in y .

$P(d)$ is the average power over all possible thresholds in $[-\infty, +\infty]$ and is the mean of the relative distribution (Handcock & Morris, 1999). Values of $P(d)$, for selected values of d and D , obtained by numerical integration, are plotted in Figure 3. The plots indicate a strongly linear relationship between the D score and $P(d)$ across the entire range of values, suggesting that, for overlapping Gaussians, D has a linear relationship with the dominance probability across the entire range of values. The non-linearities are evident at higher values of Cohen's d . This is a consequence of the D score and $P(d)$ being bounded by 2 and 1 respectively.

Single Subject Data

In this simulation, single subject data from a 2 x 2 experiment with 500 latencies per condition were analyzed. The population means in the four conditions were 300 ms, 400 ms, 600 ms, and 800 ms and the simulated means in the sample approached these values (Table 1). In each condition, the exponential τ was set at 30% of the overall mean of the ex-gaussian mixture and σ was set at 10% μ . For the standardized and rank measures the scaling was done in two possible ways. In the first, the normalization was based on two of the four conditions and the second used latencies from all four conditions to construct the standard deviation and ranking. For the standardized measures, the means in the latter "All" case are lower as the divisor is larger.

The mean latencies indicate the presence of two main effects that are qualified by an interaction. Latencies are longer under conditions of low alertness and they are longer in the choice task. Further, the task effect is twice as large under conditions of low alertness. The mean logarithm measure agrees with these main effects but shows a zero interaction as the task effect is 0.13 and 0.12 across high and low alertness.

$\bar{c}$ is clearly inappropriate in the "Low" and "High" scaling columns as there are only 2 conditions and the mean difference, independent of data, is constrained to be 2. In the "All"

columns, $\bar{c}$ has the same pattern as the mean latencies, indicating two main effects and one interaction. This is also the case for $\bar{d}$ and $\bar{D}$ when the scaling factor is based on all four conditions.

When $\bar{d}$ and $\bar{D}$ are computed separately within each alertness level, the mean difference across tasks is the same and the interaction vanishes. $\bar{F}$ and $\bar{G}$, closely related to $\bar{d}$ and $\bar{D}$ in the two group context, also show this pattern when the ranking of latencies is done separately within each alertness level. Unlike $\bar{d}$ and $\bar{D}$, $\bar{F}$ and $\bar{G}$ continue to exhibit zero interaction even when the ranking is done across all latencies. These measures, along with the mean log latency, show clear main effects and zero interaction.

This example demonstrates the susceptibility of the mean latency to create a spurious within subject interaction. Although the measures have the same rank order across the four conditions, divergences emerge when the interaction or *difference between contrasts* is considered. Due to fine scaling, the mean latency differences are much larger in the low alertness condition. This interaction disappears in the case of the logarithm, rank measures $\bar{F}$ and $\bar{G}$, and the standardized measures $\bar{d}$ and $\bar{D}$ when they are scaled within alertness levels. Fine scaling, the correlation between mean and standard deviation, is eliminated by these measures.

Response Latency Simulation

Response latencies were simulated for subjects in two contrasted conditions such that the data conformed to the following constraints: there was a positive correlation between the two mean latencies, across subjects. This was implemented by assigning a base latency to each subject and then defining condition means as deviations from base latency. The data was coarsely scaled so that the larger mean latency also had greater variability across subjects. Fine scaling, the positive correlation between the mean and standard deviation of within subject latencies, emerged as a consequence of using the ex-gaussian to generate individual latency distributions. For these data, contrast measures were evaluated for their effect sizes, internal consistencies, and correlation with base latency, an index of general information processing speed.

Preparation of Simulation Data

Each subject was assigned a base latency β from a normal distribution with $\mu = 600$, $\sigma = 100$. Base latency reflects general processing speed and mean latencies in the two conditions were positive and negative deviations from base latency. The mean latency $\bar{x}$ is smaller than the base latency by the factor f_x , $\bar{x} = \beta(1 - f_x)$. The mean latency $\bar{y}$ is larger than the base latency by the factor f_y , $\bar{y} = \beta(1 + f_y)$. f_x and f_y were uniformly sampled from the ranges $[0, 0.25]$ and $[0, 0.5]$ respectively so that $\bar{y}$ had a greater standard deviation than $\bar{x}$.

For a subject with a base latency of 700 and effect factors 0.1 and 0.3, the mean latencies were 630 ms and 910 ms. Once the subject mean latencies were assigned, the trial latencies were generated using the ex-gaussian distribution. For each subject, θ was sampled from $[0.2, 0.3]$ and was the contribution of the exponential component to the mixture. Thus $\tau_x = \theta \bar{x}$, $\tau_y = \theta \bar{y}$, $\mu_x = (1 - \theta) \bar{x}$, and $\mu_y = (1 - \theta) \bar{y}$. η was sampled from the range $[0.1, 0.2]$ and used to set the gaussian standard deviations, $\sigma_x = \eta \mu_x$ and $\sigma_y = \eta \mu_y$. Data for 500 subjects, with 100 trial latencies in each of two conditions, generated a total of 100000 trial latencies.

β , f_x , f_y , θ , and η define the contrasting ex-gaussian distributions and these, in turn, were used to generate trial latencies. The base latency β constrains $\bar{x}$ and $\bar{y}$ to be positively correlated across subjects. θ and η , along with the ex-gaussian ensure the correlation between the mean and variance of latency. As f_y is more variable than f_x , $\bar{y}$, the mean of the larger latency, is more variable than $\bar{x}$ across subjects.

Simulation Results

Table 2 and Table 3 display data generation parameters and sample statistics respectively. β , f_x , f_y , θ , and η were close to the expected values as were the parameters of the ex-gaussian. Within subjects, on average, s_y was 1.46 times larger than s_x and, between subjects,

$\lambda = \frac{s_y}{s_x} = 1.58$. The computation of $\bar{c}_\Delta$ depends on the presence of three or more conditions.

We create a dummy factor that distinguished between odd and even trials and computed the standard deviation of mean latency across four conditions to provide the denominator of $\bar{c}_\Delta$. In effect, the design has a significant main effect, a secondary factor that has zero effect and there is a zero interaction. $\bar{D}_\Delta$, $\bar{F}_\Delta$, and $\bar{G}_\Delta$ had notably higher effect sizes than $\bar{L}_\Delta$, $\log \bar{L}_\Delta$, and $\bar{d}_\Delta$.

However the effect size of $\bar{c}_\Delta$ was enormous at 7.72 and it can be attributed to the unusually low variance of the measure. The internal consistencies, based on split-half correlations across odd and even trials, seen in Table 3, are, with one exception, over .90. The exception is $\bar{c}_\Delta$ which has a negative correlation between the odd and even trials implying that it is not an empirically or even logically consistent measure. In these data the standard deviation between the four means, the scaling factor for $\bar{c}_\Delta$, correlated .999 with the IAT effect. This is not surprising because the second factor and the interaction had zero effects. Thus the extraordinarily high effect size of $\bar{c}_\Delta$ is simultaneously associated with its lack of internal coherence. The internal consistencies of the other measures are inflated by the consistency of the simulation process. In empirical data there is considerably more noise from trial to trial which along with practice, sequence, order, and interaction effects work to reduce internal consistency.

Fine scaling is seen in the correlations between s_x and $\bar{x}$ ($r=.772$) and s_y and $\bar{y}$ ($r=.812$). $\bar{L}_\Delta$ is uncorrelated with $\bar{x}$ but correlates .8 with $\bar{y}$, a consequence of coarse scaling (Sriram et al., 2006). $\bar{L}_\Delta$ correlates .32 with the base latency β and .54 with the average mean latency. $\log \bar{L}_\Delta$, $\bar{d}_\Delta$, $\bar{D}_\Delta$, $\bar{F}_\Delta$, and $\bar{G}_\Delta$ are uncorrelated with base latency and only weakly related to the average latency (see Table 4). The correlation between $\bar{L}_\Delta$ and $\log \bar{L}_\Delta$ is .928 and ranges from .87 to .88 with the other measures. $\bar{D}_\Delta$ is very strongly correlated with $\bar{F}_\Delta$ and $\bar{G}_\Delta$ ($r=.989$) and the latter are correlationally nearly identical. $\bar{c}_\Delta$ is uncorrelated with base latency and average latency but has weak relations to the other contrasts due to its complete lack of internal consistency.

Although these simulation data fall well short of extreme realism, compared to $\overline{L}_\Delta$, the scale invariant measures, in particular $\overline{D}_\Delta$, $\overline{F}_\Delta$, and $\overline{G}_\Delta$, more powerfully detected the main effect and were independent of general processing speed.

Empirical Data Analysis

The Implicit Association Test or IAT contrasts response latencies from two critical conditions (Greenwald, McGhee & Schwartz, 1998). Each IAT condition is composed of 60 trials, in consecutive 20 trial and 40 trial blocks. Participants receive classification training prior to these critical blocks. Feedback is provided on incorrect responses and participants are obliged to make the correct response in order to advance to the next trial.

We examined data from the Age IAT, which concerns preference for *Young* over *Old*. In the Age IAT, participants first learn to distinguish “Good” from “Bad” words and “Young” from “Old” faces in separate blocks. These two classifications are combined in two critical conditions. In the *GoodYoung* condition pictures of young people and “Good” words are associated with a common response. The alternative response is shared by pictures of old people and “Bad” words. In the *GoodOld* condition pictures of old people and “Good” words are associated with a common response and the alternative response is shared by pictures of young people and “Bad” words. Preference for *Young* over *Old* is reflected by shorter latencies in *GoodYoung* and longer latencies in *GoodOld*. In the vast majority of participants, the age IAT revealed a clear preference for *Young* over *Old*. Data were obtained from visitors to the Internet site <http://implicit.harvard.edu> during the final quarter of 2006 (for more general information about the methodology and empirical results relating to these data see Nosek, 2005 and Nosek et al., 2006).

The Age IAT

Design. The data were analyzed using a 2 x 2 x 2 x 2 repeated measures ANOVA. The first factor measures the IAT effect and is labeled *Congruity* (I). The second factor, *Practice* (P), distinguishes the first block from the second block. The third factor, *Type* (T) separates responses to word stimuli from responses to picture stimuli. As trials were randomized within each block, approximately half the trials necessitate a shift from the response on the previous trial and approximately half the trials involve pressing the same response key as the preceding trial. The fourth factor, *Shifting* (S), makes a distinction between non-shift and shift trials. The first trial in each block was considered as a shift trial although omitting this from the analysis does not change the results.

Besides the Congruity effect resulting from faster *GoodYoung* responses, main effects of *Practice* and *Shifting* were expected. Participants were expected to speed up on the second block and to be faster on non-shift trials as they do not need to initiate a response shift. *Type* effects were not predicted and are not utilized in evaluating comparisons of effect sizes across measures.

Scaling procedure. Standardization for *c*, *d*, and *D* was based on latencies from all 16 conditions. These measures divide the untransformed latencies by, the standard deviation across the 16 condition mean latencies, the average within condition standard deviation, and the standard deviation of all trial latencies, respectively. The *F* and *G* transformations were based on ranking within subject latencies across 16 conditions. Participants with more than 25% errors, 4.5% of the sample, were excluded from analysis. 5407 participants, with complete data in the 16 conditions, were used as the basis for the Age IAT analysis.

Effect size comparisons. As the sample size is very large, significance testing recedes in importance and the focus is on effect size. On 11 of the 15 possible effects, at least one measure had effect size, $\eta_p^2 > 0.01$; the F ratio at this threshold is over 50. These effect sizes are listed in Table 5. Focusing on the substantial main effects of I, P, and S, it is clear that all the scale invariant measures had larger effects than the untransformed latencies. The maximum variability across the measures is seen for *Shifting* where the logarithm and standardized measures have larger effects than raw latency, and the rank measures *F* and *G* have still stronger effects.

Similarity among measures. The three substantial main effects were used to gauge the similarity among the measures. The main effects, collapsed across other conditions, were computed and the correlations, across subjects, are presented in Tables 6a, 6b and 6c. $\overline{F}_\Delta$ and $\overline{G}_\Delta$ are correlationally nearly identical, with *r*'s over .99. Likewise, $\overline{d}_\Delta$ and $\overline{D}_\Delta$ are very similar, with *r*'s greater than .98. For all three effects, $\overline{L}_\Delta$ is closest to $\overline{\log L}_\Delta$ which, however, has higher correlations with the other scale invariant measures, in particular $\overline{G}_\Delta$. Among the standardized measures, $\overline{c}_\Delta$ is distinct but is most similar to $\overline{D}_\Delta$. As the number of treatment conditions increases, the between condition variance will dominate most of the total variance and the correlation between $\overline{c}_\Delta$ and $\overline{D}_\Delta$ will approach 1.

As the pattern of effect sizes in Table 5 suggest, $\overline{L}_\Delta$ and the rank measures are most weakly related for *Shifting* and have about 60% distinct variance. This factor most clearly exemplifies a case in which $\overline{L}_\Delta$ is relatively small but the distributional separation, as captured by the rank contrasts, is more consistent than indicated by the $\overline{L}_\Delta$ effect size. Although it is outside the scope of this article, the relative magnitude of interaction effects across $\overline{L}_\Delta$ and the rank measures exhibit sharp discrepancies.

Internal consistency of contrasts. For each of the I, P, and S main effects, eight contrasts were defined based on combinations of factors other than the main effect under consideration. The contrasts were internally consistent in that these eight parallel measures correlated positively with each other. Cronbach α 's, listed in Table 7, indicate that *Practice* has the least internal consistency whereas *IAT* and *Shifting* have much greater consistency.

For all three effects, the standardized contrast $\overline{c}_\Delta$ has the lowest internal consistencies, a trend that appeared in an extreme fashion in the response latency simulation. Scaling by the standard deviation of the condition mean latencies for $\overline{c}_\Delta$ results in lower internal consistencies, even as it has high effect sizes. This problem is particular acute when the number of conditions is small, e.g., 4, and most of the variance is due to one factor, as in the response latency simulation.

Relation to general processing speed. One of the desirable attributes of latency contrasts, articulated time and again (Faust et al. 1999, Greenwald et al., 2003), is that it not be statistically constrained, like $\overline{L}_\Delta$ (Sriram et al., 2006), to be associated with general processing speed. Two plausible measures of general speed were derived, the average latency in the two double classification critical conditions (*Critical RT*, $M=865$, $SD=232$) and the average latency during the single classification training blocks (*Training RT*, $M=715$, $SD=151$). In these data they

correlated .792, a high degree of overlap but retained considerable unique variance. Table 8 lists the regression coefficients predicting main effects from average critical and training RTs.

Given the large sample size, standardized regression coefficients greater than .08 can be considered significant as they exceed the threshold of $p < .001$. For all three effects, $\overline{L}_A$ shares considerable variance with the average critical and training RTs. $\overline{\log L}_A$ reduces this but retains considerable variance with general speed indices. The standardized and rank measures have 3 to 10% shared variance with general speed and, achieving considerable independence from measures of general speed. However they do share some non-zero variance with general speed that may be theoretically interpretable.

Homogeneity of effect variances. It is curious that the rank contrasts, based on the uniform rank distribution have superior power to $\overline{L}_A$ on parametric ANOVA that assumes normally distributed variables. $\overline{F}_A$ is based on averaging fractional ranks; these rapidly converges to normality, due to the central limit theorem, as highly skewed outliers in the raw latencies are eliminated by ranking (Hall & Wang, 2004); in the case of $\overline{G}_A$, convergence is even faster as the overall within subject distribution is gaussian. In addition to normality, independence, and homogeneity of variance, the repeated measures design requires homogeneity of treatment difference variances (Maxwell and Delaney, 1990, pp. 471-474). If the variances of these differences are homogeneous then the *sphericity* assumption of repeated measures ANOVA is satisfied. The problem of unequal variances in the case of reaction time measures has been recognized and non-parametric tests have been proposed (Zumbo & Colombe, 1997).

Table 9 displays SDs of the IAT effects across 8 cells in the design. $\overline{L}_A$ violates sphericity as the SD of the effects are heterogeneous, particularly across the two blocks. The logarithm and the standardized measures do reduce this heterogeneity but there is still quite a large variability across blocks. The rank contrasts display the most homogeneous effects across the various conditions and best satisfy ANOVA assumptions in these data. Truncation of outliers often leads to loss of valid information in condition contrasts (Ulrich & Miller, 1994). $\overline{F}_A$ and $\overline{G}_A$ are extremely robust to outlier data and can incorporate them in the analysis with minimal side effects.

Between subjects effects. Within subject ranking and standardization loses information about group differences. That is, all between group main effects and interaction terms involving only group factors are set to zero. A variant of G that retains between groups differences can be constructed in two steps. In the first step, all latencies across all subjects are ranked in order and gaussian latencies, with mean=500 and standard deviation=100, for each fractional rank is computed. This first round of ranking is done at the level of raw latencies across subjects. Next a second round of latency ranking, within subjects, is conducted. These ranks are converted to a within-subject gaussian with a standard deviation of 100. However the mean of this within-subject gaussian is constrained to be the subject mean of the gaussian latencies computed in the first round of ranking. The second ranking is functionally identical to that used in earlier data analyses; however, instead of using 500 as the mean, the within subject gaussian is centered around the average gaussian rank for the subject obtained from the first round of ranking.

When latency data are transformed in this manner, both within and between subjects factors can be assessed using ANOVA. Three such between subject factors were identified for these data, based on responses to a demographics questionnaire. 4744 subjects, a subset of the

data analyzed earlier, provided their gender (1596 males, 3148 females) and their age (1361 were 18 or younger, 2027 ranged from 19 to 30, and 1356 subjects were 31 years or older). They performed the IAT in one of two counterbalanced orders, with either *YoungGood* preceding *OldGood* (2376 subjects) or *OldGood* preceding *YoungGood* (2368 subjects). The effects provided in Table 5 were not influenced to any notable degree by focusing on this sample subset and our attention will be on the between subjects factors, *Age* (A), *Sex* (X), and *Order* (O).

Table 10 displays effects involving at least one between subjects factor that have $\eta_p^2 > .01$ for one of the three contrasts under consideration. Only 11 of 112 effects exceeded the .01 threshold. The main effects of age (older subjects were slower) and gender (women subjects were faster) were powerfully detected by $\overline{G}_\Delta$ as were effects involving *Order* and *Shifting*. Men showed stronger *IAT* effects and this is most powerfully detected by $\overline{G}_\Delta$. The interaction between subject Age and the IAT effect is particularly strong for $\overline{L}_\Delta$ but much weaker for $\overline{G}_\Delta$ and $\log \overline{L}_\Delta$, suggesting that a large part of this interaction in $\overline{L}_\Delta$ is an artifact of scaling.

Interpreting gaussian latency. The gaussian latency is a relative measure because it is based on rank ordering latencies, both between and within subjects. In the above example, the standard deviation for both distributions was set at 100 and the value of the grand mean latency, across all subjects was set at 500. This choice of values eases interpretation because a mean difference in gaussian latency units simultaneously encodes effect size information. A difference of 25 gaussian ms within subjects can be interpreted as 0.25 SDs of the within subject latency distribution and likewise for between subjects effects. In choosing this or any other transformation, there is unavoidable loss of information. Percentile differences have a behavioral interpretation as the probability of latencies from one condition dominating those from contrasting conditions. In the case of multiple conditions and groups, as in repeated measures ANOVA, the average percentile for any particular condition reflects its position, with respect to the entire distribution of latencies across conditions.

The results in this article have implications for the manner in which latency contrasts are conceptualized. The rank order information in latencies captures a considerable portion of the reliable variance in effects. Standardization in the two condition case is similar to ranking although standardization clearly diverges from ranking in analyzing contrasts among a number of conditions within a single subject. Although the gaussian rank is very closely related to the fractional rank, unlike F , G weighs deviations that are away from the mean more heavily than those that are near the mean and this distinction increases statistical power. The gaussian rank latency is a normally distributed latent variable underlying response latencies. The simplicity and utility of this conceptualization needs to be evaluated thoroughly on a broad range of data to better appreciate its potential and its limitations.

Potential Applications

Bialystok, Craik, Klein, and Viswanathan (2004) reported that bilinguals outperformed monolinguals in the Simon task. The superiority of bilinguals was larger in older participants. The Simon effect contrasts a slower, more variable incongruent condition with a faster, less variable congruent condition. Consequently, $\overline{L}_\Delta$ of the Simon effect is expected to correlate with general processing speed. Compared to younger subjects, older subjects have larger latencies that are more variable. The Age x Bilingualism interaction in the Simon effect should be checked using scale invariant measures.

Laver and Burke (1993) conducted a meta-analysis of the finding that semantic priming effects are larger for older adults than for young adults. Priming effects in individual experiments are small. Other investigators report the interaction of semantic priming with word frequency, reading skill, and extent of visual degradation (Madden, 1988). Larger semantic priming $\overline{L}_A$ was evident in conditions (low frequency words, poorly skilled readers and degraded stimuli) that are associated with increased latency means and variances. Which interactions might be reversed, weakened, or strengthened using scale invariant measures? If poorly skilled readers also have longer and more variable latencies, it is possible that the larger priming $\overline{L}_A$ associated with them may be exaggerated.

In a meta-analysis of the interaction of the Stroop effect with age, Verhaeghen and Meersman (1998) conclude that there were no significant age effects despite large overall differences in Stroop effects across young and old subjects (271 vs. 511 ms). Reanalysis of raw data using scale invariant measures may establish a moderate age interaction effect. Spieler et al. (1996) observe that old adults make many more extremely slow responses in the interference condition. Using $\overline{G}_A$, such responses could result in a larger Stroop effect for adults without increasing the error variance associated with such outliers.

Scale invariant measures may have a role in contexts other than response latency. Scaling evoked response potentials has been controversial. McCarthy and Wood (1985) proposed a form of scaling to eliminate the possibility of spurious Location x Condition interactions caused by variability in strength of the source of the voltage potential across the scalp. This article became influential in promoting the practice of normalizing scalp amplitudes prior to further analysis. More recently, the McCarthy and Wood argument has been sharply criticized for making unrealistic assumptions in the simulation that first established the existence of the purported problem of spurious interactions (Haig, Gordon, & Hook, 1997). The current recommendation is to discontinue the widespread practice of amplitude normalization of ERP data (Urbach & Kutas, 2002). There may be a role for scale invariant measures in the assessment of ERP main effects and interactions, particularly if the ERP mean across locations, time slices, conditions, and subjects is associated with the ERP variance.

Conclusion

We propose scale invariance and validity maximization as defining properties of admissible latency contrast measures. In the absence of scale invariance, comparisons of latency contrasts, between subjects and within subject, become difficult to understand and hard to communicate. Given multiple scale invariant measures, the ones that maximize construct validity in a domain may be preferred on pragmatic grounds. Unlike mean latency differences, these measures were not spuriously related to processing speed and were able to detect valid effects more readily. $\overline{G}_A$, in particular, appears to have the simultaneous advantages of (i) being readily interpretable (ii) detecting both between and within subjects effects (iii) reducing correlational biases, (iv) addressing fine scaling at the level of single subject data, (v) homogenizing effect variance to better satisfy ANOVA assumptions, and (vi) exhibiting superior internal consistency and external validity. These advantages may not be limited to the data analyzed in this report but may indicate that $\overline{G}_A$ has the potential to improve the analysis of latency contrasts generally.

Scale invariance and validity maximization are objectives at different levels of analysis. Scale invariance is assessed *deductively* by logical and mathematical analysis of the numerical properties of a measurement procedure. Validity maximization is assessed *inductively* by empirical examination of a measure's statistical reliability, its power to detect associations with causally related constructs and exhibit independence from causally unrelated constructs. The solution for the Behrens-Fisher problem (Ghosh & Kim, 2001), concerning inference for the mean difference between two normal populations whose ratio of variances is unknown, is predicated on the assumption that the mean difference validly represents the construct being measured. While mean differences are suited to many contexts, $\overline{L}_\Delta$ is a flawed measure of response latency contrasts. In comparison to $\overline{L}_\Delta$, the recommended scale invariant measures are statistically powerful, reduce correlational biases, and exhibit superior construct validity.

References

- Balota, D.A., & Chumbley, J.I. (1984). Are lexical decisions a good measure of lexical access? The role of word frequency in the neglected decision stage, *JEP: Human Perception and Performance*, 10, 340-357.
- Bartlett, M. S. (1947). The use of transformations. *Biometrics*, 3, 39-52.
- Bialystok, E., Craik, F. I. M., Klein, R., & Viswanathan, M. (2004). Bilingualism, aging, and cognitive control. *Psychology and Aging*, 19, 290-303.
- Bush, L. K., Hess, U., & Wolford (1993). Transformations for within-subject designs: A Monte Carlo investigation. *Psychological Bulletin*, 113, 566-579.
- Carpenter, R. H. S. (1999). A neural mechanism that randomises behaviour. *Journal of Consciousness Studies*, 1, 13-22.
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences*. LEA: Hillsdale, NJ.
- Cliff, N. (1996). *Ordinal methods for behavioral data analysis*. LEA: Mahwah, NJ.
- Faust, M. E., Balota, D. A., Spieler, D. H., & Ferraro (1999). Individual differences in information processing and amount: Implications for group differences in response latency. *Psychological Bulletin*, 125, 777-799.
- Fazio, R.H., Sanbonmatsu, D.M., Powell, M.C., & Kardes, F.R.(1986). On the automatic activation of attitudes, *Journal of Personality and Social Psychology*, 50, 229-238.
- Fisher, D. L. & Glaser, R. A. (1996). Molar and latent models of cognitive slowing: Implications for aging, dementia, depression, development, and intelligence. *Psychonomic Bulletin and Review*, 3, 458-480.
- Friedman, M. (1937). The use of ranks to avoid the assumption of normality implicit in the analysis of variance. *Journal of the American Statistical Association*, 32, 675.
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology*, 74, 1464-1480.
- Greenwald, A. G., Nosek, B. A., Banaji, M. R. (2003). Understanding and using the Implicit Association Test: I. An improved scoring algorithm. *Journal of Personality and Social Psychology*, 85, 197-216.
- Ghosh, M., & Kim, Y-H. (2001). The Behrens-Fisher problem revisited: A Bayes-Frequentist synthesis. *Canadian Journal of Statistics*, 29, 5-17.

- Hall, P., & Wang, Q. (2004). Exact convergence rate and leading term in central limit theorem for Student's t statistic. *Annals of Probability*, 32, 1419-1437.
- Handcock, M. S., & Morris, M. (1999). *Relative Distribution Methods in the Social Sciences*. Springer-Verlag: New York.
- Haig, A. R., Gordon, E., & Hook, S. (1997). To scale or not to scale: McCarthy and Wood revisited. *Electroencephalography and Clinical Neurophysiology*, 103, 323-325.
- Heathcote, A., Brown, S., & Mewhort, D. J. K. (2002) Quantile maximum likelihood estimation of response time distributions. *Psychonomic Bulletin and Review*, 9, 394-401.
- Laver, G. D., & Burke, D. M. (1993). Why do semantic priming effects increase in old age? A meta-analysis. *Psychology and Aging*, 8, 34-43.
- Luce, R. D. (1986). *Response times, their role in inferring elementary mental organization*. New York: Oxford University Press.
- MacLeod, C. M. (1991). Half a century of research on the Stroop effect: An integrative review. *Psychological Bulletin*, 109, 163-203.
- Madden, D. J. (1988). Adult age differences in the effects of sentence context and stimulus degradation during visual word recognition. *Psychology and Aging*, 3, 167-172.
- Maxwell, S. E., & Delaney, H. D. (1990). *Designing experiments and analyzing data: A model comparison perspective*. Wadsworth: Belmont, CA.
- McCarthy, G. & Wood, C. C. (1985). Scalp distribution of event-related potentials: An ambiguity associated with analysis of variance models, *Electroencephalography and Clinical Neurophysiology*, 62, 203-208.
- Meiran N. (1996). Reconfiguration of processing mode prior to task performance. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 22 1423-1442.
- Meyer, D. E., Osman, A. M., Irwin, D. E., & Yantis, S. (1988). Modern Mental Chronometry. *Biological Psychology*, 26, 3-67.
- Miller, J. (1988). A warning about median reaction time. *Journal of Experimental Psychology: Human Perception and Performance*, 14, 539-543.
- Myerson, J., Adams, D. R., Hale, S., & Jenkins, L. (2003). Analysis of group differences in processing speed: Brinley plots, Q-Q plots, and other conspiracies. *Psychonomic Bulletin & Review*, 10, 224-237.
- Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner & G. W. Humphreys (Eds), *Basic processes in*

- reading: *Visual word recognition* (pp. 264-336). Hillsdale, NJ, England: Lawrence Erlbaum Associates, Inc.
- Nosek, B. A. (2005). Moderators of the relationship between implicit and explicit evaluation. *Journal of Experimental Psychology: General*, 134, 565-584.
- Nosek, B. A., Banaji, M. R., & Greenwald, A. G. (2002). Harvesting implicit group attitudes and beliefs from a demonstration website. *Group Dynamics*, 6, 101-115.
- Nosek, B. A., Greenwald, A. G., & Banaji, M. R. (2006). The Implicit Association Test at age 7: A methodological and conceptual review. In J. A. Bargh (Ed.), *Social Psychology and the Unconscious: The Automaticity of Higher Mental Processes* (pp. 265-292). Psychology Press.
- Nosek, B. A., Smyth, F. L., Hansen, J. J., Devos, T., Lindner, N. M., Ranganath, K. A., Smith, C. T., Olson, K. R., Chugh, D., Greenwald, A. G., & Banaji, M. R. (2006). Pervasiveness and Correlates of Implicit Attitudes and Stereotypes. Unpublished manuscript.
- Patel, J. C. Kapadia, H. & Owen, D. B. (1976). *Handbook of Statistical Distributions*. New York: Marcel Dekker, Inc.
- Proctor, R. W. & Cho, Y. S. (2006). Polarity correspondence: A general principle for performance of speeded binary classification tasks. *Psychological Bulletin*, 3, 416-442.
- Ratcliff, R., Gomez, P., & McKoon, G. (2004). A diffusion model of the lexical decision task. *Psychological Review*, 1, 159-182.
- Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, 9, 347-356.
- Salthouse, T. A. (1985). Speed of behavior and its implication for cognition. In J.E. Birren & K. W. Schaie (Eds.), *Handbook of the psychology of aging* (2nd ed., pp. 400-426. New York: Van Nostrand Reinhold.
- Salthouse, T. A., & Hedden, T. (2002). Interpreting reaction time measures in between-group comparisons. *Journal of Clinical and Experimental Neuropsychology*, 24, 858-872.
- Spieler, D. H. (2001). Modelling age-related changes in information processing. *European Journal of Cognitive Psychology*, 13, 217-234.
- Snedecor, G. W., & Cochran, W. G. (1989). Statistical methods (8th ed). Ames (IA): Iowa State University Press.
- Spieler, D. H., Balota, D. A., & Faust, M. E. (1996). Stroop performance in healthy younger and older adults and in individuals with dementia of the Alzheimer's type. *Journal of Experimental Psychology: Human Perception and Performance*, 22, 461-479.

- Sriram, N., Greenwald, A. G., Nosek, B. A. (2006). Correlational biases in mean response latency differences. Unpublished manuscript. University of Virginia.
- Ulrich, R. & Miller, J. (1994). Effect of truncation on reaction time analysis. *Journal of Experimental Psychology: General*, 123, 34-80.
- Urbach, T. P. & Kutas, M. (2002). The intractability of scaling scalp distributions to infer neuroelectric sources, *Psychophysiology*, 39, 791-808.
- Verhaeghen, P., & Meersman, L. D. (1998). Aging and the Stroop effect: A meta analysis. *Psychology and Aging*, 13, 120-126.
- Verhaeghen, P., Cerella, J., Semenec, S. C., Leo, M. A., Bopp, K. L., & Steitz, D. W. (2002). Cognitive efficiency modes in old age: Performance on sequential and coordinative verbal and visuospatial Tasks. *Psychology and Aging*, 17, 558-570.
- Wagenmakers, E-J., Grasman, R. P. P. P., Molenaar, P. C. M. (2005). On the relation between the mean and variance of a diffusion model response time distribution. *Journal of Mathematical Psychology*, 49, 195-204.
- Zumbo, B. D. & Coulombe, D. (1997). Investigation of the robust rank-order test for non-normal populations with unequal variances: The case of reaction time. *Canadian Journal of Experimental Psychology*, 51, 139-149.

TABLE 1
Simple RT vs. Choice RT

<i>Measure</i>	<i>High Alertness</i>				<i>Low Alertness</i>			
	<i>Simple</i>		<i>Choice</i>		<i>Simple</i>		<i>Choice</i>	
$\bar{L}$	302		407		602		803	
	(96)		(126)		(190)		(252)	
$\overline{\log L}$	2.46		2.59		2.76		2.89	
	(.121)		(.116)		(.119)		(.118)	
<i>Scaling</i>	<i>High</i>	<i>All</i>	<i>High</i>	<i>All</i>	<i>Low</i>	<i>All</i>	<i>Low</i>	<i>All</i>
$\bar{c}$	5.80	1.58	7.80	2.12	6.01	3.15	8.01	4.19
$\bar{d}$	2.72	1.82	3.66	2.45	2.72	3.63	3.63	4.83
$\bar{D}$	2.45	1.16	3.29	1.56	2.46	2.31	3.28	3.08
$\bar{F}$	.353	.189	.648	.370	.356	.632	.645	.810
$\bar{G}$	450	392	550	464	451	541	549	602

TABLE 2

Data Generation Parameters

<i>Parameter</i>	<i>M</i>	<i>SD</i>
β	610	97
f_x	.119	.070
f_y	.254	.145
θ	.250	.029
η	.149	.029
μ_x	403	71
μ_y	574	114
σ_x	60	15
σ_y	85	23
τ_x	134	29
τ_y	191	44

TABLE 3

Effect Size and Internal Consistency

<i>Statistic</i>	<i>M</i>	<i>SD</i>	<i>Cohen's d</i>	<i>Cronbach α</i>
$\overline{x}$	537	95		
$\overline{y}$	764	150		
$\overline{L_{\Delta}}$	228	111	2.05	.945
$\overline{\log L_{\Delta}}$	.152	.064	2.36	.943
$\overline{c_{\Delta}}$	1.93	.250	7.72	-.835
$\overline{d_{\Delta}}$	1.28	.562	2.28	.909
$\overline{D_{\Delta}}$	1.00	.342	2.94	.918
$\overline{F_{\Delta}}$	.325	.106	3.07	.931
$\overline{G_{\Delta}}$	107.4	33.7	3.19	.930

TABLE 4

Correlations among Base Latency, Average Latency, and Contrasts

	β	$\overline{x+y}$	$\overline{L_{\Delta}}$	$\overline{\log L_{\Delta}}$	$\overline{c_{\Delta}}$	$\overline{d_{\Delta}}$	$\overline{D_{\Delta}}$	$\overline{F_{\Delta}}$
$\overline{x+y}$	.893							
$\overline{L_{\Delta}}$	.320	.542						
$\overline{\log L_{\Delta}}$	.008	.223	.928					
$\overline{c_{\Delta}}$	-.078	.006	.360	.401				
$\overline{d_{\Delta}}$	.011	.208	.879	.943	.396			
$\overline{D_{\Delta}}$	-.005	.198	.871	.948	.479	.978		
$\overline{F_{\Delta}}$	.003	.209	.880	.951	.484	.960	.989	
$\overline{G_{\Delta}}$	-.002	.202	.874	.945	.503	.954	.989	.998

TABLE 5

Age IAT: Repeated Measures ANOVA Effect Sizes (η_p^2)

<i>Effect</i>	<i>L</i>	<i>Log L</i>	<i>C</i>	<i>d</i>	<i>D</i>	<i>F</i>	<i>G</i>
<i>I</i>	.414	.516	.567	.512	.539	.557	.558
<i>P</i>	.373	.474	.459	.463	.480	.508	.519
<i>S</i>	.193	.296	.283	.253	.260	.319	.361
<i>T</i>	.023	.014	.021	.018	.020	.010	.021
<i>IP</i>	.087	.057	.065	.100	.096	.010	.012
<i>IS</i>	.004	.012	.003	.005	.005	.009	.016
<i>IT</i>	.017	.057	.053	.048	.050	.075	.076
<i>PS</i>	.058	.055	.077	.070	.072	.044	.038
<i>PT</i>	.082	.100	.101	.096	.100	.084	.088
<i>TS</i>	.160	.151	.195	.190	.187	.075	.079
<i>PTS</i>	.051	.040	.055	.056	.056	.015	.025

TABLE 6a
Age IAT Effect Correlations

	$\overline{L_{\Delta}}$	$\overline{\log L_{\Delta}}$	$\overline{c_{\Delta}}$	$\overline{d_{\Delta}}$	$\overline{D_{\Delta}}$	$\overline{F_{\Delta}}$
$\overline{\log L_{\Delta}}$	.934					
$\overline{C_{\Delta}}$	.735	.849				
$\overline{d_{\Delta}}$	.820	.917	.894			
$\overline{D_{\Delta}}$	.808	.911	.934	.982		
$\overline{F_{\Delta}}$	.790	.933	.877	.913	.915	
$\overline{G_{\Delta}}$	.800	.936	.891	.927	.932	.995

TABLE 6b
Practice Effect Correlations

	$\overline{L_{\Delta}}$	$\overline{\log L_{\Delta}}$	$\overline{c_{\Delta}}$	$\overline{d_{\Delta}}$	$\overline{D_{\Delta}}$	$\overline{F_{\Delta}}$
$\overline{\log L_{\Delta}}$	.923					
$\overline{C_{\Delta}}$	.744	.876				
$\overline{d_{\Delta}}$	.822	.920	.902			
$\overline{D_{\Delta}}$	.813	.921	.937	.984		
$\overline{F_{\Delta}}$	.737	.910	.888	.884	.894	
$\overline{G_{\Delta}}$	.762	.926	.906	.912	.922	.992

TABLE 6c
Shifting Effect Correlations

	$\overline{L_{\Delta}}$	$\overline{\log L_{\Delta}}$	$\overline{c_{\Delta}}$	$\overline{d_{\Delta}}$	$\overline{D_{\Delta}}$	$\overline{F_{\Delta}}$
$\overline{\log L_{\Delta}}$	.893					
$\overline{C_{\Delta}}$	.688	.837				
$\overline{d_{\Delta}}$	.798	.902	.887			
$\overline{D_{\Delta}}$	.781	.901	.925	.981		
$\overline{F_{\Delta}}$	.623	.869	.823	.824	.835	
$\overline{G_{\Delta}}$	.673	.900	.860	.872	.884	.988

TABLE 7
Internal Consistency of IAT, Practice, and Shifting (Cronbach α)

	$\overline{L_{\Delta}}$	$\overline{\log L_{\Delta}}$	$\overline{c_{\Delta}}$	$\overline{d_{\Delta}}$	$\overline{D_{\Delta}}$	$\overline{F_{\Delta}}$	$\overline{G_{\Delta}}$
Age IAT	.774	.821	.628	.748	.706	.815	.800
Practice	.679	.687	.497	.575	.560	.666	.660

Shifting .809 .807 .704 .744 .741 .796 .790

TABLE 8
Effects Predicted by Average RT on Critical and Training Blocks

	<i>Age IAT</i>			<i>Practice</i>			<i>Shifting</i>		
	β_{crit}	β_{train}	R^2	β_{crit}	β_{train}	R^2	β_{crit}	β_{train}	R^2
$\overline{L_\Delta}$	.654	-.227	.243	.390	.195	.300	.228	.127	.251
$\log \overline{L_\Delta}$	.548	-.256	.138	.216	.223	.173	.098	.227	.096
$\overline{c_\Delta}$	.255	-.191	.023	-.038	.313	.080	-.214	.331	.042
$\overline{d_\Delta}$	.316	-.203	.037	.046	.250	.083	-.100	.284	.045
$\overline{D_\Delta}$	.288	-.186	.031	.015	.278	.084	-.137	.306	.045
$\overline{F_\Delta}$	.365	-.185	.058	-.005	.312	.095	-.134	.273	.029
$\overline{G_\Delta}$	.351	-.179	.054	.001	.327	.094	-.161	.279	.032

TABLE 9

Means and SDs of the Age IAT effect

	<i>Block 1</i>				<i>Block 2</i>			
	<i>Words</i>		<i>Pictures</i>		<i>Words</i>		<i>Pictures</i>	
	<i>No Shift</i>	<i>Shift</i>	<i>No Shift</i>	<i>Shift</i>	<i>No Shift</i>	<i>Shift</i>	<i>No Shift</i>	<i>Shift</i>
$\overline{L}_\Delta$	279 (.468)	246 (.435)	238 (.472)	228 (.526)	182 (.273)	159 (.256)	136 (.266)	147 (.259)
$\overline{\log L}_\Delta$	.110 (.150)	.099 (.133)	.093 (.151)	.082 (.134)	.085 (.107)	.076 (.097)	.064 (.105)	.066 (.092)
$\overline{c}_\Delta$	1.01 (1.44)	.912 (1.33)	.785 (1.40)	.755 (1.45)	.656 (.984)	.495 (.990)	.386 (.682)	.419 (.662)
$\overline{d}_\Delta$	.830 (1.19)	.737 (1.07)	.658 (1.15)	.635 (1.21)	.556 (.757)	.489 (.683)	.386 (.682)	.419 (.661)
$\overline{D}_\Delta$	.623 (.884)	.553 (.800)	.491 (.868)	.471 (.899)	.420 (.562)	.370 (.513)	.289 (.513)	.315 (.493)
$\overline{F}_\Delta$	.184 (.250)	.169 (.214)	.157 (.247)	.130 (.189)	.164 (.194)	.155 (.177)	.126 (.194)	.132 (.165)
$\overline{G}_\Delta$	63.6 (86.8)	55.3 (71.1)	52.2 (83.4)	43.9 (66.0)	55.9 (67.5)	50.0 (58.5)	41.6 (65.9)	43.2 (55.2)

TABLE 10

Age IAT: Between Subjects Effects (η_p^2)

<i>Effect</i>	<i>L</i>	<i>Log L</i>	<i>G</i>
<i>A</i>	.094	.117	.135
<i>X</i>	.008	.010	.010
<i>IA</i>	.011	.002	.003
<i>IO</i>	.006	.005	.014
<i>IX</i>	.015	.017	.019
<i>TA</i>	.003	.017	.028
<i>TX</i>	.012	.019	.019
<i>IPO</i>	.047	.050	.052
<i>ISO</i>	.002	.001	.012
<i>ITAO</i>	.009	.011	.009
<i>ITSO</i>	.013	.012	.009

Figure Captions

Figure 1. Brinley Plot of Young and Older Adult Mean RTs for Two Sets of Tasks (from Spieler, 2001).

Figure 2. Fine Scaling in Individual Latency Distributions.

Figure 3. Cohen's d , D Score, and Dominance Probability.