# Explicitly modeling the effects of aging on response time

ROGER RATCLIFF Northwestern University, Evanston, Illinois

DANIEL SPIELER Stanford University, Stanford, California

and

GAIL McKOON Northwestern University, Evanston, Illinois

Research into the effects of aging on response time has focused on Brinley plots. Brinley plots are constructed by plotting mean response times for older subjects against those for young subjects for a set of experimental conditions. The typical result is a straight line with a slope greater than 1 and a negative intercept. This linear function has been interpreted as showing that aging leads to a general slowing of cognitive processes. In this article, we show that the slope of the Brinley plot is actually a measure of the relative standard deviations of older versus young subjects' response times; it is not a measure of general slowing. We examine current models of the effects of aging on mean response time and show how they might be reinterpreted. We also show how a more comprehensive model, Ratcliff's diffusion model (1978), can account for Brinley plot regularities and, at the same time, provide an account of accuracy rates, the shapes of response time distributions, and the relative speeds of error and correct response times, aspects of the data about which models designed to account for Brinley plots are mute. We conclude by endorsing a research approach that applies explicit models to response time data in aging in order to use the parameters of the model to interpret the effects of aging.

As people age, their response times in cognitive tasks slow. The degree to which their responses slow is often graphed as a Brinley plot (Brinley, 1965; Salthouse & Somberg, 1982). A Brinley plot is constructed either by computing the mean response times for older subjects and young subjects separately for each condition of an experiment and then plotting the means against each other or by plotting the mean response times for individual subjects against each other, with the fastest older subject's mean response time plotted against the fastest young subject's mean response time, the next fastest older subject's mean response time against the next fastest young subject's mean response time, and so on. Either way they are constructed, Brinley plots show a striking regularity: Across a wide range of response times and types of experiments, the function is always approximately a straight line with a slope of around 1.5 (ranging from 1 to 3.0). In other words, response times for older subjects can be predicted from response times for young subjects by a simple linear transformation.

The traditional avenue to interpretation and understanding of the Brinley regularity is to take the linear function as the central datum to be explained, and efforts are directed toward understanding how response time and cognition in general change with age. We review two models of this type: a linear model proposed by Cerella (1985) and the information loss model proposed by Myerson, Hale, Wagstaff, Poon, and Smith (1990). Both models instantiate the effect of aging on response time as a general slowing effect: All the components of processing in an experimental task are assumed to be slower for older people than for young people, by the constant factor that is the slope of the Brinley plot.

Although it has been argued recently that the regularities in Brinley plots (their linearity and the similarity of the values of their slopes) provide only weak constraints on possible models of the effects of aging on response time (Cerella, 1994; Fisk & Fisher, 1994; Myerson, Wagstaff, & Hale, 1994; Perfect, 1994), the Brinley plot still serves as a benchmark, in that experimental data are expected to exhibit linear Brinley functions and theories of cognitive processing are expected to produce them.

In this article, we present a new insight into Brinley plots. We point out that Brinley plots are quantile–quantile plots (Q–Q plots), which are well studied in statistics (see Chambers, Cleveland, Kleiner, & Tukey, 1983; Thomas & Ross, 1980) and for which there is a substantial body of theory. Looking at Brinley plots as Q–Q plots leads to a new interpretation of their linearity and their slope. In brief, the value of the slope is determined by the ratio of

This research was supported by National Institute of Mental Health Grant HD MH44640 and National Institute for Deafness and other Communication Disorders Grant R01-DC01240. Correspondence concerning this article should be addressed to R. Ratcliff, Department of Psychology, Northwestern University, Evanston, IL 60208 (e-mail: r-ratcliff@nwu.edu).

the standard deviation of older subjects' response times to the standard deviation of young subjects' response times. The slope cannot be interpreted as a general slowing factor. The feature of a Brinley plot that captures the fact that older subjects are typically slower than young subjects is the function's intercept.

Reinterpretation of the Brinley plot requires a new approach to modeling. The key finding that a model must explain is not that older subjects are slowed by a constant factor, relative to young subjects, but rather that the distribution of response times for older subjects is wider. We argue that, in order to accomplish this, a model must identify the component processes that are involved in a cognitive task and then show how one or another of those processes can be affected by aging in such a way as to produce the appropriate distributions of response times. In the later sections of this article, we use simulations to illustrate how one model. Ratcliff's diffusion model for simple twochoice decisions (Ratcliff, 1978; Ratcliff & Rouder, 1998, 2000; Ratcliff, Van Zandt, & McKoon, 1999), can do this. An important conclusion that follows from the simulations is that Brinley plots provide only weak constraints on theory. Any theory that appropriately scales the difficulty of an experimental task onto response times (such as the diffusion model or competing models; e.g., P. L. Smith & Vickers, 1988) will produce a Brinley plot with appropriate characteristics. A much fuller consideration of response time data—including, for example, accuracy rates, error response times, and the shapes of response time distributions-will be needed to constrain models.

We begin below with an explanation of Q-Q theory and how it applies to Brinley plots. Then we review the earlier models (Cerella, 1985; Myerson et al., 1990) and their reinterpretations in terms of Q-Q theory. Last, we illustrate the variety of ways in which the diffusion model, a model that identifies components of processing, can simulate Brinley plot data.

### **QUANTILE-QUANTILE PLOTS**

There is one major mystery about Brinley plots: Why are they almost always linear? There also is a second puzzling regularity, that the intercept of the Brinley plot is negative when the slope of the Brinley plot is greater than 1 and the larger the slope, the more negative is the intercept. The theory of Q-Q plots provides a framework for understanding these regularities. Consider the distributions of mean response times from the conditions of an experiment, one distribution of condition means for older subjects and one for young subjects. What Q-Q theory shows is that when the condition means for the older and the young subjects are plotted against each other, the slope of the function is the ratio of the standard deviations of the distributions, and when the distributions have the same shape, the function is automatically linear. Also, from the Q-Q analysis, it is to be expected that the intercept is negative and becomes more negative as the slope

increases. Thus, the issue of explaining the shape of the Brinley plot is translated into an issue of explaining the shapes of the distributions of mean response times. It is the shapes of these distributions that should be the target of models of processing: If a model produces similarly shaped distributions of response times for older and younger subjects across conditions, the model will automatically produce linear Brinley plots.

## **Linear Brinley Plots**

Brinley plots are most often constructed by plotting the mean response time for older subjects in each condition of an experiment against the mean response time for young subjects in each condition. The plot then has as many points as there are conditions in the experiment (conditions in an experiment may be different values of word frequency, stimulus onset asynchrony, list length, study time, stimulus discriminability, etc.). The values of mean response times form a distribution across conditions, one distribution for older subjects' means and one for young subjects' means.

Given that the mean for an experimental condition is a point from a distribution of means, each mean is actually a *quantile* of the distribution. The quantiles of a distribution are the points that divide the total frequency in the distribution into parts. For example, the median point divides the distribution into halves (.5 quantiles), the three quartile points divide the distribution into quarters (.25 quantiles), and so on. A Q–Q plot simply plots the quantile points of another distribution. So, for example, the score at the .1 quantile for the other distribution, and so on. The means plotted on a Brinley function are quantiles for the distribution are a plotted on a Brinley function are quantiles for the distribution are a plotted.

If the distributions on which a Q–Q plot is based are normal, the Q–Q plot is a straight line. The slope of the O–O plot is the ratio of the standard deviations of the distributions, and the intercept is a function of the means of the two distributions and their standard deviations. To understand why this is so, consider, first, two normal distributions, each with a mean of 0 and a standard deviation of 1; that is, each is the standard normal distribution. Plotting quantiles from these two distributions against each other [e.g., 10%, Q(.1) = -1.28; 20%, Q(.2) = -0.84; 30%, Q(.3) = -0.52; etc.] gives a straight line [in standard z tables, this would correspond to  $Q(\alpha) = z(\alpha)$ ]. If the response time scores of older and young subjects come from normal distributions, their quantiles are the quantiles of the standard N(0,1) normal [e.g., Q(.1) = -1.28] multiplied by the standard deviation of their distribution ( $\sigma_{o}$ for the older subjects and  $\sigma_v$  for the young subjects) and added to the mean of their distribution ( $\mu_{o}$  and  $\mu_{v}$ , respectively). The quantiles in our response time setting represent different levels of difficulty of the experimental task (x), and so the transformations from the quantiles for the standard normal to the older subjects' distribution is  $Q_0 =$ 

 $\sigma_0 x + \mu_0$ , and the transformation to the young subjects' distribution is  $Q_y = \sigma_y x + \mu_y$ . Eliminating difficulty x, we obtain a straight line with slope =  $\sigma_0 / \sigma_y$  and intercept =  $\mu_0 - \mu_y \sigma_0 / \sigma_y$ . The full equation for a point on the Q–Q (Brinley) plot is

$$Q_{o} = (\sigma_{o}/\sigma_{v})Q_{v} + \mu_{o} - \mu_{v}\sigma_{o}/\sigma_{v}.$$
 (1)

These results and the equation hold if the two distributions from which the Q–Q plots are derived are not normal but, rather, logistic, Cauchy, gamma (with a fixed "number" parameter), exponential, Weibull (with fixed exponent), uniform, or ex-Gaussian (with the parameter of the exponential component being a constant times the standard deviation of the normal component; Ratcliff & Murdock, 1976). These distributions all have the form that their quantiles can be derived from the standard form by multiplying by a constant and adding a constant. Thomas and Ross (1980) provide a detailed discussion of the properties that give rise to linearity for these distributions, and Chambers et al. (1983, chap. 6) provide a more introductory discussion of Q–Q plots, along with a list of the properties of Q–Q plots for several of these distributions.

The important point is that, if the Q–Q plot for a set of conditions in an experiment is produced from two distributions with approximately the same shape, the slope is the ratio of the standard deviation of the two distributions (as long as the range of the points on the plot is significantly greater than the variability in each point; see the Appendix). In other words, any model that produces distributions of older and young subjects' mean response times so that the distributions have the same shape will automatically produce a linear Q–Q (Brinley) plot with a slope equal to the ratio of the standard deviations.

The equation does not hold if the distribution of older subjects' means is more positively skewed than the distribution of younger subjects' means (e.g., an exponential for the older subjects and a normal for young subjects)that is, more positively skewed than would be obtained from two different members of the same distribution family (e.g., two exponentials with different means). In this case, the Q-Q plot is positively accelerated (e.g., an exponential plotted against a normal distribution). That is, the function is U-shaped, with the left and right tails of the function lying above a straight line through the central part of the function. This shape has occasionally been obtained empirically (Cerella, 1990; Nebes & Madden, 1988; see also Myerson et al., 1990). Even though the equation does not hold, the slope of the straight line fitted to the function turns out, in many cases, to be the ratio of standard deviations, as in Equation 1.

All of the points just made about the relationship between Q–Q plots and Brinley plots also hold if the Brinley plot is constructed, not by plotting the mean response times of older versus young subjects for each condition of an experiment, but by plotting means for individual older and young subjects against each other. The important point, that the slope of a Brinley plot is the ratio of standard deviations of older and young response times, is a fundamental reinterpretation of the Brinley plot: It means that the slope of the Brinley plot shows *nothing* about slowing of older subjects relative to young subjects. What the slope does show is that there is greater variability across conditions in the mean response times of older subjects than in the mean response times of young subjects. Older subjects *are* typically slower than young subjects, but it is the intercept, which is a function of mean response times and their standard deviations, that shows the slowing.

Figure 1 illustrates this interpretation of the Brinley plot. The top left panel shows two normal distributions of response times, one for young subjects with a mean of 600 and a standard deviation of 100 [N(600, 100)] and one for older subjects with a mean of 700 and a standard deviation of 150 [N(700,150)]. These distributions give typical response times for older versus young subjects. Plotting randomly generated scores from these distributions in a Brinley plot (i.e., a Q-Q plot) shows the typical 1.5 slope that has been interpreted as showing a general slowing. But the slope actually represents the ratio of the standard deviations of the two distributions (150/100). The right-hand panels show the same slope obtained from two distributions in which the older subjects are, on average, faster than the young subjects. The distribution for young subjects is N(600, 100), and the distribution for older subjects is N(500, 150). The slope is still 1.5, but the older subjects are faster. The points made by these example distributions apply whether the response times are taken to be the response times of individual subjects or the mean response times for experimental conditions.

Figure 2 illustrates the point mentioned above, that the slope of the Brinley plot is still the ratio of standard deviations even when the response times are sampled randomly from nonnormal distributions. The top row of panels shows the Brinley plots, the second row shows the distributions for young subjects, and the third row shows the distributions for older subjects. The three distributions have quite different shapes, and yet, in each case, the ratio of the older to young standard deviations is within 1% of the value of the Brinley slope.

## The Relation Between a

## **Brinley Plot's Slope and Its Intercept**

Another regularity that has emerged in the Brinley plot literature is that, when the slopes and intercepts of the Brinley plots from many experiments are plotted against each other, the result is a straight line, as is shown in Figure 3. Cerella (1991; see also Maylor & Rabbitt, 1994) and Faust, Balota, Spieler, and Ferraro (1999) provided tables of the Brinley intercepts and slopes from a number of experiments. The intercepts varied between +214 msec and -693 msec, and the slopes varied between 0.96 and 2.63. We plotted each experiment's intercept



Figure 1. Examples of the translation from distributions of mean response times across subjects to Brinley or quantile–quantile (Q–Q) plots.

against its slope (following Cerella, 1991) to produce the straight-line function with an intercept of 572 msec and a slope of -441 msec that is shown in Figure 3.

The fact that this function is a straight line follows directly from the O–O analysis, under several reasonable assumptions about how aging might affect response times. To show this, we begin with normal distributions of older and young subjects' response times and examine a variety of possible relations between the two distributions (the distributions could be either distributions of condition means or distributions of individual subjects' means). We fixed the young subjects' distribution to have a mean of 600 msec and a standard deviation of 100 msec, varied the older subjects' mean from 600 to 750 msec, as is shown in Table 1, and varied the older subjects' standard deviation from 100 to 250 msec, as is shown in the first column of the table. Then, for each pair of young and older standard deviations, we computed the slope of the Brinley plot (from Equation 1), and from each pair of young and older standard deviations and each pair of young and older means, we computed the intercept of the Brinley plot (again, from Equation 1). The resulting Brinley slopes and intercepts are shown in the table.

Before proceeding, two things are worth pointing out about the values in Table 1. First, as the slope of the Brinley plot increases, the intercept decreases, just as in Figure 3. Also, the intercept is mostly negative when the slope is greater than 1. Only a relatively large mean response time, coupled with a relatively small standard deviation for the older subjects, produces a positive intercept.

The main question is whether the values in Table 1 produce the linear function in Figure 3. We show that they do, in three cases that span a range of reasonable assumptions about how older subjects' response times vary with respect to young subjects'. In each case, the value of the slope is negative, and the intercept is positive. (The values in the table can also be used to evaluate other assumptions about how older subjects' response times vary against young subjects' response times.)

For the first example, we assume that older and young subjects have the same mean, 600 msec, and that only the standard deviation for older subjects varies, from 100 to 250 msec. The resulting Brinley slopes are given in the second row of the table, and the Brinley intercepts in the third row. Plotting these slopes and intercepts against each other gives a straight line with a slope of -600 msec and an intercept of 600 msec.

Second, we assume that both the mean for the older subjects' distribution and their standard deviation increase relative to the young subjects, and for purposes of illustration, we assume that they increase by the same amount, 30 msec, beginning from a standard deviation of 100 msec and a mean of 600 msec. The resulting Brinley slopes are again the values in the second column of the table, and the Brinley intercepts are the values on the diagonal from top left to bottom right. Plotting these slopes and inter-



Figure 2. Simulated shapes of Brinley (Q–Q) plots derived from exponential, normal, and uniform distributions of mean response times across subjects. The distributions in the bottom row have larger mean and larger standard deviation (1.5 times larger) than the distributions in the middle row. The Brinley functions plot the distributions in the middle row on the *x*-axis and the distributions in the bottom row on the *y*-axis.

cepts against each other, the function is a straight line with a slope of -500 msec and an intercept of 500 msec.

Third, we assume that the mean for older subjects is larger than the mean for young subjects by a constant amount (90 msec larger for purposes of illustration) and that the older subjects' standard deviation varies from 100 to 250 msec. The Brinley slopes are again the values in the second column of the table, and the Brinley intercepts are the values in the "690" column of the table. The plot of these slopes and intercepts is again a straight line, with a slope of -466 msec and an intercept of 597 msec (very near the values for the line in Figure 3).

These three cases illustrate three ways older subjects' response times might vary relative to young subjects' response times. In each case, the plot of the intercepts of the Brinley plot versus the slopes of the Brinley plot gives a negative slope and a positive intercept, with the third case being very close to the observed values in Figure 3. Thus,



Figure 3. Plots of intercepts versus slopes of Brinley plots for 39 experiments from Cerella (1985) and Faust, Balota, Spieler, and Ferraro (1999).

from the Q–Q analysis, it follows that any reasonable assumption about how older and young response times differ will produce the observed relationship between the slopes and intercepts of the Brinley plots from different experiments.

#### Summary

Brinley plots are Q–Q plots, and this fact gives us insight into multiple aspects of the plots.

1. Brinley plots are linear, because the distribution of mean response times across experimental conditions for older subjects has about the same shape as the distribution of means across conditions for young subjects. If, as has sometimes been found, the Brinley function is slightly accelerated upward, the distribution of older subjects' means has greater skew (i.e., an extra large number of slow response times) than does the distribution of young subjects' means.

2. The slope of a Brinley plot is greater than 1, because older subjects' mean response times have greater variability than young subjects' means, *not* because of general slowing.

3. The intercept of a Brinley plot is negative, because of the relative placements of the means and standard deviations of the older and young subjects' distributions of mean response times; older subjects' means are larger than younger subjects' means, but not by too much, relative to their standard deviations.

Table 1
Slopes and Intercepts of Brinley Plots
for Normal Distributions of Response Time

Standard Deviation in Response Times for Simulated Older	Slope of	Mean Response Time for Simulated Older Subjects (msec)						
Subjects (msec)	Brinley Plot	600	630	660	690	720	750	
100	1.0	0	30	60	90	120	150	
130	1.3	-180	-150	-120	-90	-60	-30	
160	1.6	-360	-330	-300	-270	-240	-210	
190	1.9	-540	-510	-480	-450	-420	-390	
220	2.2	-720	-690	-660	-630	-600	-570	
250	2.5	-900	-870	-840	-810	-780	-750	

Note—Mean response time for young subjects is 600 msec; standard deviation in response time for young subjects is 100 msec. 4. The slopes and intercepts of Brinley functions are negatively correlated (as is shown in Figure 3), because when the standard deviation of older subjects' means increases relative to the standard deviation of the young subjects' means, the slope of the Brinley plot increases and the intercept decreases (irrespective of whether the older subjects' means increase along with their standard deviation), since the ratio of standard deviations appears in both the slope and the intercept of Equation 1.

A strong argument has been made in the literature about the lack of a need to model task performance in order to understand the effects of aging on processing time. The argument goes as follows: Because Brinley functions show that the effect of aging is a general and linear slowing and because this slowing appears across a range of tasks, from tasks that can be performed very quickly to tasks that are very slow, hypotheses about the reason for general slowing are all that are needed to explain the effects of aging.

However, the analysis in terms of Q–Q plots provides a different way of viewing the Brinley plot regularities namely, that they arise from a larger spread in the distribution of mean response times for older subjects than for young subjects. This pattern could occur for many reasons (e.g., older subjects may adopt more conservative response criteria, or the information they extract from a stimulus may be of lower quality than that for young subjects), which we illustrate in the next sections of this article, and the reasons could be different in different tasks. Distinguishing among the different explanations for the larger spread of older subjects' distributions will require examination not just of the mean response times plotted in Brinley functions, but also of other aspects of the response time data, such as accuracy, error response times, and the shapes of response time distributions.

## MODELS FOR THE EFFECTS OF AGING ON MEAN RESPONSE TIMES

Most of the theoretical work in aging and response time has focused on the general slowing hypothesis, the idea that the aging effect is a general slowing effect. The assumption has been that the Brinley plot shows this slowing, with the slope of the plot giving the factor by which older subjects slow, relative to young subjects. However, the Q–Q reinterpretation of the Brinley plot slope indicates that the aging effect is not a general slowing effect as measured by the slope of the Brinley function. The models *can* be reinterpreted to be consistent with the Q–Q analysis, but at the expense of losing their general slowing account of aging.

#### Cerella's (1985) Linear Model

Cerella (1985) proposed that response times for older subjects ( $RT_o$ ) are a linear function of response times for young subjects ( $RT_y$ )—that is,  $RT_o = a RT_y + b$ . We describe his proposal in the slightly generalized form in which response time is a function of task difficulty (as in Cerella, 1994, Figure 1). Figure 4 (top left panel) shows mean response time plotted against difficulty for four values of a, represented by the lines labeled "1," "2," "3," and "4," with the intercept (b) the same for the four functions. Like Cerella's function, this generalization gives the Brinley pattern, because response time is a linear function of task difficulty:

$$RT = ax + b, \tag{2}$$

where *x* represents difficulty and *b* represents the floor on response time. Then,  $RT_o = a_o x + b_o$ , and  $RT_y = a_y x + b_y$ , and with simple algebra to eliminate *x*,

$$RT_{o} = (a_{o}/a_{v})RT_{v} + b_{o} - b_{v}a_{o}/a_{v}.$$
 (3)

In other words, response time for older subjects is a linear function of response time for young subjects (a Brinley plot), with a slope of  $a_o/a_y$  and an intercept of  $b_o - b_y(a_o/a_y)$ . In the top left panel of Figure 4, a group of young subjects might produce, for example, the line marked with 2s for response time as task difficulty increased. The line for a group of older subjects would be higher—for example, the line marked with 3s (the lines have the same intercept, b, but if different intercepts were assumed, the analysis would carry through in the same way with the same qualitative conclusions).

The findings just outlined for the assumption of a linear relationship between difficulty and response time generalize to other functions that scale difficulty into response time:

$$RT = a f(x) + b, \tag{4}$$

where f(x) is some transformation of difficulty, the same transformation for older and young subjects, and where older and young subjects differ only in parameters a and b. The top left panel of Figure 4 shows response time as a linear function of difficulty. For response time as an exponential function of difficulty—for example, RT = a $\exp(x/c) + b$ —the scaling functions show the same general pattern (Figure 4, bottom left panel), and the same derivations as those for the linear model carry through (because elimination of the term  $\exp(x/c)$  from the versions of this equation for young and older subjects gives Equation 3). Thus, the form of the function relating difficulty to response time does not have to be linear; rather, it has to be of the form f(x), where x is difficulty, so that when f(x) is eliminated from the equations for young and older subjects' response times, the resulting relationship between young and older subjects' response times is linear.

The expressions for the slope and intercept for Cerella's (1985) model and our generalizations (Equation 3) look very much like those for a Q–Q plot (Equation 1); *a* and *b* in Equation 1 correspond to  $\sigma$  and  $\mu$ , respectively, in Equation 3. The Q–Q analysis is more general than Cerella's model and accommodates nonlinear Brinley functions. Here, we show how to relate the Q–Q analysis and Cerella's model by using three different examples. First, suppose that the conditions of an experiment are equally



Figure 4. Plots of mean response time against number of stages of difficulty or drift rate for four response time models: A linear model (Cerella, 1985); the Myerson, Hale, Wagstaff, Poon, and Smith (1990) information loss model; an exponential model; and the diffusion model (Ratcliff, 1978).

spaced as a function of difficulty. Then, for Cerella's model, with a linear transformation from difficulty to response time, response times for older and young subjects will both have uniform distributions (e.g., the rightmost panels of the middle and bottom rows of the panels in Figure 2), and the Brinley function will be linear (as in the top row, rightmost panel of Figure 2). For the Q-Q analysis, with the distributions of response time uniform (both having the same shape), the Q-Q (Brinley) plot will be linear, with a slope that is the ratio of the standard deviations of the distributions. Second, suppose that the relationship between the difficulty of the conditions of an experiment and response time is exponential (e.g., RT = a $\exp(x/c) + b$ , leftmost panels in the middle and bottom rows of Figure 2) rather than linear and that conditions are equally spaced as a function of difficulty. Then, the Brinley plot from a Cerella-type model will be linear, but with a greater density of points in the bottom left corner of the function than in the top right corner (Figure 2, top row, leftmost panel). For the Q–Q analysis, the distribu-

tions of response times would be exponential (leftmost panels of the middle and bottom rows in Figure 2), and the slope of the O–O function would be the ratio of their standard deviations. Third, the distributions of response times across conditions could have a nonuniform distribution with a linear relationship between difficulty and response time. This could arise if the range of difficulty is not sampled uniformly-for example, the .1, .2, .4, and .8 quantiles in the range of difficulty might be sampled, rather than there being uniform sampling. This would lead to a distribution of points with more points from easy conditions than from difficult conditions. Then, the distributions of condition means might look like the distributions shown in the leftmost panels of the middle and bottom rows in Figure 2, and the resulting Brinley plot would be like that in the top row, leftmost panel. This example would be indistinguishable (in Brinley plot analyses) from the case in which the mapping from difficulty to response time was exponential (the second point above). In sum, for any reasonable ways of sampling, values from the difficulty dimension (with distributions for older and younger subjects from the same families) will lead to response times that are consistent with Cerella's model.

However, there is one limitation: If the function relating difficulty to response time is not the same for older and young subjects—that is, if the response times for older subjects come from a different kind of distribution than the young subjects' response times (e.g., older subjects' from an exponential distribution and young subjects' from a uniform distribution)—Brinley plots will be accelerated upward (if the old distribution is more skewed than the young distribution), and Cerella's (1985) linear model will not fit the data. The Q–Q analysis, on the other hand, is still consistent with an upward accelerated function.

## Myerson et al.'s (1990) Information Loss Model

Myerson et al.'s (1990) model describes performance on cognitive tasks as resulting from a series of processing stages, where the stages are not further specified, so that the model can apply across a variety of tasks. Total response time is assumed to be the sum of the processing times for each stage,  $RT = \sum T_k$ , and the time for each stage k is inversely related to the amount of information available at that stage:  $T_k = D/I_k$ , where D is a constant. The information at stage k is  $I_k = I(1-p)^k$ , where I and p are constants, p determining the amount of information lost at each stage. This model is called an information loss model, because  $I_k$  decreases as the number of stages increases. With the assumption that aging principally affects the parameter p, this model approximates linear Brinley plots well (the model produces approximately a power function), at least as well as the linear model (cf. Cerella, 1985).

The top right panel of Figure 4 shows mean response time as a function of number of stages in the model for four values of p. Just as with the linear and exponential functions, the effect of aging is a scaling effect: Increasing the number of processing stages required by a task increases response time more for older subjects than it does for young subjects, because the value of p is larger for older subjects.

For Myerson et al.'s (1990) model, if either a uniform or a nonuniform distribution of values of difficulty were selected in an experiment, the distribution of mean response times across conditions for older subjects would be slightly more skewed than that for young subjects, and the Brinley (Q–Q) plot would be accelerated upward slightly. The slope of the linear approximation to this function would be nearly the same as the ratio of standard deviations of the distributions. Data with this nonlinear function would be interpreted as different shaped distributions across conditions or subjects for young and older subjects in a Q–Q analysis.

A third model, similar to Myerson et al.'s (1990), should be mentioned, and that is Cerella's (1990) neural network. The idea is that cognitive tasks are performed by neural networks. It is similar to Myerson et al.'s in that processing in a task is assumed to be composed of a series of stages. The stages are abstract and not tied to the specific task being performed. Aging produces breaks in the links (stages) of the network, and these breaks require the creation of new links (stages), which results in longer neural pathways among the nodes in the network. These longer pathways lead to longer response times. The probability of breaking a link is assumed to be an exponential function of age. With auxiliary assumptions, this model can account for several different patterns of the data showing aging effects on response times. An analysis similar to that for the Myerson et al. model can be carried out to relate this model to Q-Q plots.

#### Summary

The Cerella (1985) and the Myerson et al. (1990) models translate (scale) task difficulty into response times for young and older subjects. The models can be reinterpreted into Q–Q plot terms, and for Cerella's model, the mathematical equation relating response times for older and young subjects can even be mapped directly onto the Q–Q equation. But the reinterpretation is a radical one: The original idea of the models, that aging is a general slowing reflected in the slope of the Brinley plot, has to be given up and replaced with the understanding that what the slope actually shows is that older subjects differ from young subjects in terms of the standard deviations of their response time distributions across conditions.

## SCALING ISSUES

Scaling is at the heart of the issue of how subjects slow with age. The linear, exponential, and stage models illustrated in Figure 4, all general slowing models, show scaling: Response times for older subjects are scaled as a function of task difficulty, so that each next level of difficulty (or number of stages) slows them more than it does the response times for young subjects. In other words, the young subjects follow one function (e.g., the "1" function), and older subjects follow another function (e.g., the "3" function).

However, so far we have not considered a critical aspect of performance: interactions between speed and accuracy, and how speed-accuracy criteria might differ between older and young subjects. If speed-accuracy criteria do differ between older and young subjects, the increasing differences in response times as a function of task difficulty might be due to older subjects' setting more conservative criteria. Consider what young subjects typically do when faced with moving from an experimental condition of less difficulty to one of greater difficulty: They slow their responses but they do not slow them enough to maintain their same level of accuracy; instead, they allow their accuracy to decrease. Older subjects might try not to allow their accuracy to decrease and, in doing so, might slow more than the young subjects. This idea, that older subjects sacrifice speed for gains in accuracy, has been suggested frequently (e.g., Basowitz & Korchin, 1957; Silverman, 1963; Strayer & Kramer, 1994; Thorndike, Bregman, Tilton, & Woodyard, 1928). In fact, for some tasks, older adults may make fewer errors than young adults (e.g., Stroop and flanker tasks; Spieler, Balota, & Faust, 1996, in press). What we stress here is that the ability to trade speed for accuracy might explain why the difference in response times between two conditions is greater for older than for young subjects; according to this explanation, increasing task difficulty does not affect older subjects more than young subjects in any way other than in the accuracy criteria they set.

The scaling that results from speed versus accuracy instructions was demonstrated in an experiment by Ratcliff and Rouder (1998, Experiment 1), using young subjects. Subjects were asked to judge whether squares presented on a PC screen were *bright* or *dark*. The squares were made up of  $64 \times 64$  pixels—some pixels black, some white, and some a neutral gray—and the bright squares had more white pixels than did the dark squares. Hard conditions were those in which the numbers of white and black pixels were about the same; easy conditions were those in which there were many more pixels of one kind than of the other. When subjects were instructed to respond as fast as possible, overall, their correct response times averaged about 350 msec, and the difference between hard and easy conditions averaged less than 50 msec. But when they were instructed to respond as accurately as possible, their correct response times averaged about 700 msec, and the difference between hard and easy conditions averaged about 200 msec. In other words, when response times were fast overall, the difference between conditions was dramatically smaller than when response times were slow overall. What is important to appreciate is that this scaling is not related to the scaling by difficulty that is shown for the linear, exponential, and stage models in Figure 4. For the subjects in Ratcliff and Rouder's (1998) experiment, the difficulty of the bright/ dark decision was not greater at short response times than at longer ones; rather, the reason that the difference between easy and hard conditions was larger at long than at short response times was the scaling of response times onto different points on a speed-accuracy tradeoff function. We could plot a Brinley function for the young subjects' response times in the accuracy conditions against their response times in the speed conditions; the slope would be 2.82 and the intercept would be -282 msec, exactly in the range of typical Brinley values for older versus young subjects, but the slope would have nothing to do with differences in difficulty. (The largest differences in accuracy between the speed and the accuracy conditions in Ratcliff and Rouder's (1998) Experiment 1 were only about 5%. This means that small differences in accuracy between older and young subjects cannot be taken to mean that there are no speed-accuracy criterion differences between them.) This exercise shows that differences in speed-accuracy criteria settings could be largely responsible for the typical Brinley plots of older versus young subjects' response times.

The Ratcliff and Rouder (1998) data also illustrate another important point about response times—namely, the nonlinearity of response time differences (cf. Pachella, 1974). At relatively short response times, the variance in response times is small, and the differences among experimental conditions are small. Differences among conditions are smallest when response time approaches floor, where the floor for a given task and an individual subject is the fastest response time at which accuracy begins to rise above chance in response signal paradigms (see, e.g., Dosher, 1984; Ratcliff, 1978; Reed, 1973, 1976; Wickelgren, 1977). At longer response times, both the variance and the differences among experimental conditions become larger.

The points just elaborated show the complexity involved in response time scaling. Greater differences among experimental conditions for older than for young subjects may be due to speed-accuracy tradeoff differences, to the greater effect of difficulty on older subjects' responses, or to some mixture of these. Also, there might be other factors, such as older and young subjects' operating at different levels of difficulty on the difficulty-response time function, differences in encoding and response output times, increased variability in processing, or differences in the likelihood of subjects' rejecting the first decision on some proportion of trials and engaging in a second decision process. Current models like Cerella's (1985) and Myerson et al.'s (1990) offer no way to sort out the different factors. There is no way, for example, to use the models to evaluate the possibility of speed-accuracy tradeoffs. What is needed is a model that elucidates the various possibilities. We believe that any such model will necessarily be a model the lays out the components of processing in the tasks to which it is applied. In the next section, we show how the diffusion model accomplishes these goals.

## THE DIFFUSION MODEL

Ratcliff's diffusion model (Ratcliff, 1978, 1981, 1985, 1988; Ratcliff & Rouder, 1998, 2000; Ratcliff et al., 1999) is designed to account for response time and accuracy in experimental paradigms in which subjects are asked to make two-choice decisions. The heart of the model is a decision process in which two components of processing are separated: The information from the stimulus is separated from the response criteria according to which a decision is made. The model has been successful in dealing with data from a range of experimental paradigms but has not yet been applied to the effects of aging on response time. Here, we show how this could be done.

The diffusion model assumes that decisions are made by a noisy process that accumulates information over time toward one of two response criteria (or boundaries; see Figure 5, top panel). When one of the criteria is reached, a response is initiated. The mean rate of accumulation of information is called the *drift rate*, and it is determined by the quality of the information driving the decision process.



Figure 5. Top panel: An illustration of the diffusion model showing sample paths from a discrete approximation (i.e., small steps instead of a continuous process). The dotted lines represent the case in which criteria are moved close to the starting point so that the three paths terminate at the points marked T. Bottom panel: An illustration of scaling in the diffusion model, with two different boundary positions  $(a_1 \text{ and } a_2)$  and three different drift rates  $(v, v + \Delta, \text{ and } v + 2\Delta)$ .

For example, in recognition memory, a well-remembered item would have a large value of drift rate. In a visual perception task, an item that was easy to see would have a larger drift rate than an item that was not so easy to see. As is shown in Figure 5, top panel, there is variability around the mean as information is accumulated. The figure shows three processes, each with the same (positive) drift rate. Variability causes one of them to reach the upper boundary quickly, one to reach the upper boundary more slowly, and the third to reach the lower boundary in error. If the boundaries were moved nearer the starting point (the dotted horizontal lines), the processes would terminate at the points marked T—that is, two of them would end in errors. Moving the boundary positions is the way the model accounts for speed–accuracy tradeoffs: The closer the boundaries to the starting point, the faster are responses, and the more likely are errors. The model produces accurate quantitative predictions for mean response times for correct responses and error responses. The model



Figure 6. Examples of how normal distributions of drift rate for older and young subjects map into skewed distributions of response time, with greater standard deviations for older than for young subjects.

also automatically and accurately generates positively skewed response time distributions (Ratcliff, 1978; Ratcliff & Rouder, 1998; Ratcliff et al., 1999).

In the past, response time models had considerable difficulty in dealing with the relative speeds of correct versus error responses. Empirically, response times for errors are sometimes longer than response times for correct responses, sometimes shorter, and sometimes the relationship between error and correct response times varies for an individual subject across the conditions of an experiment (see Ratcliff & Rouder, 1998; Ratcliff et al., 1999; P. L. Smith & Vickers, 1988). The diffusion model can account for these varying patterns. Errors slower than correct responses are caused by the drift rate's varying not only within a trial (as is shown in Figure 5, top panel), but also across trials. In other words, it is assumed that ostensibly equivalent stimuli are not always encoded equally well from one trial to another. Without across-trial variability, the model (like any random walk or diffusion model with constant mean drift rate across trials) predicts that error response times are the same as correct response times. For error responses faster than correct responses, the diffusion model assumes variability across trials in the starting point (as was first proposed by Laming, 1968; see also Ratcliff, 1981). With these two sources of variability, the crossover patterns of error versus correct response times that appear in some data are exactly predicted by the model (errors faster than correct responses when accuracy is high, errors slower than correct responses when accuracy is low). Currently, no other model is capable of producing these patterns of results (see Ratcliff & Rouder, 1998; Ratcliff et al., 1999).

The diffusion model is meant to apply to two-choice decisions that are relatively fast and composed of a singlestage decision process—that is, to many tasks that we perform in real life many times an hour. This can be contrasted with multiple-stage decision processes that might be involved in, for example, reasoning tasks, card sorting tasks, and so forth; for such tasks, other modeling techniques are available (e.g., Fisher & Glaser, 1996). As a rule of thumb, we would not want to apply the diffusion model to experiments in which mean response times are much longer than 1-1.5 sec (although this is a rough guideline, rather than an absolute rule). With a multistage decision process, the diffusion process could be assumed to be the decision process used in each of the individual stages.

The parameters of the diffusion model correspond to the components of the decision process as follows: The starting point of the accumulation of evidence is z, the upper boundary is a, and the lower boundary is set to 0. For the simulations described in this article, the boundaries were assumed to be symmetric about the starting point, and so, z = a/2. The amount of variability in the mean drift rate across trials is assumed to be normally distributed, with a standard deviation of  $\eta$ , and the amount of variability in starting point is assumed to have a uniform distribution, with a standard deviation of  $s_z$ . For each different stimulus condition in an experiment, it is assumed that the rate of accumulation of evidence is different, and so, there is a different value of drift, v, for each condition. Finally, there is a parameter  $T_{\rm er}$  that represents the nondecisional components of response time. Within trial variability, s was kept constant in all the simulations, because

	Minimum	Maximum	Mean	., 110101110111			·	
Subject	Boundary	Boundary	Boundary	Minimum	Maximum	Mean		
Туре	Parameter (a)	Parameter $(a)$	Parameter $(a)$	Drift Rate (v)	Drift Rate $(v)$	Drift Rate (v)	Intercept (s)	Slope
Young	.14	.14	.14	.3	.7	.5	_	_
Old	.14	.14	.14	.25	.65	.45	-0.053	1.17
Old	.14	.14	.14	.2	.6	.4	-0.082	1.27
Old	.14	.14	.14	.1	.5	.3	-0.191	1.62
Old	.14	.14	.14	.15	.65	.4	-0.269	1.67
Old	.14	.14	.14	.1	.6	.35	-0.348	1.94
Old	.14	.14	.14	.0	.5	.25	-0.402	2.18
Old	.14	.14	.14	.0	.6	.3	-0.464	2.22
Young	.08	.12	.1	.3	.3	.3	_	_
Old	.12	.16	.14	.3	.3	.3	0.073	1.03
Old	.10	.16	.13	.3	.3	.3	-0.147	1.45
Old	.14	.20	.17	.3	.3	.3	-0.085	1.53
Old	.08	.16	.12	.3	.3	.3	-0.374	1.92
Old	.10	.18	.14	.3	.3	.3	-0.401	2.08
Old	.08	.18	.13	.3	.3	.3	-0.671	2.62
Old	.08	.20	.14	.3	.3	.3	-0.813	2.99

Table 2 Slopes and Intercepts of Brinley Plots From Diffusion Model Simulations

Note—The other parameters of the diffusion model were z = a/2,  $\eta = .12$ ,  $T_{er} = 0.3$  sec, and s = .1.

it is a scale parameter for the diffusion process (i.e., if it were doubled, other parameters could be multiplied or divided by two to produce exactly the same fits of the model to the data).

The components of the model that are the most likely candidates for explaining age-related differences in response times are drift rate, boundary position, and  $T_{er}$ . Older subjects might extract less information from the stimulus, and so, their drift rates might be lower than those for young subjects. They might also set boundary positions wider to make accurate performance more likely. Or they might be slower in the nondecisional components of processing. Across-trial variability might also differ between young and old, and this would have only a small effect on correct response times but would have a large effect on error response times. However, because error response times are rarely included in reports of response time data, we did not investigate the possible interactions of across-trial variability and aging in the simulations conducted for this article.

#### Illustrations of Scaling and the Diffusion Model

The diffusion model can produce the linear slope and negative intercept of a typical Brinley plot under any of a variety of possible assumptions about differences in processing between older and young subjects. We begin by illustrating this point with two examples.

First, we illustrate what happens in the diffusion model when older subjects are operating over a different part of the response time–difficulty function than young subjects. In other words, an increase in difficulty from one condition to another hurts performance for older subjects more than for young subjects, because the task is, overall, more difficult for the older subjects. Figure 5 (bottom panel) shows how difficulty, which corresponds to drift rate, maps

to response time. As difficulty increases from one condition to a second condition, older subjects' drift rates might move from  $v + \Delta$  to v, and young subjects' drift rates from  $v + 2\Delta$  to  $v + \Delta$  (i.e., the same difference  $\Delta$  in difficulty). The change in response times for older subjects  $(b_2, \text{ with boundary position } a_1)$  would be much greater than the change for young subjects  $(b_1, with boundary)$ position  $a_1$ ). The same point is illustrated in the bottom right panel of Figure 4: Older subjects might move from a drift rate of .3 in an easy condition to .2 in a hard condition, and younger subjects might move from a drift rate of .4 to .3. Assuming the same boundary positions (e.g., a = .16, the line marked with 3s) for the older and young subjects, the change in response times between the easy and the hard conditions would be greater for older than for young subjects.

The lefthand panel of Figure 6, which expands a part of the bottom right panel of Figure 4, shows the distributions of response times that would result under this assumption—that is, the assumption that an increase in difficulty from one condition to another hurts performance for older subjects more than for young subjects. The distribution of drift rates across older subjects is centered at a lower value of drift rate than is the distribution of drift rates for young subjects, and this results in the distributions of response times shown in the figure on the lefthand axis. Both distributions of response times are skewed to the right, and the older subjects' response times are generally slower, and their distribution of response times is more spread out than that for the young subjects.

Second, we illustrate what happens in the diffusion model if the drift rate is the same for older and young subjects but older subjects set more conservative decision criteria—that is, the older subjects set wider response boundaries. In Figure 5, bottom panel, drift rates might be



Figure 7. Simulated Brinley plots with drift rate alone varying between older and young subjects (top panel) or boundary positions alone changing between older and young subjects (bottom panel).

 $v + 2\Delta$  for an easy condition and  $v + \Delta$  for a hard condition for both older and young subjects. If the boundary for the young subjects was  $a_2$  and the boundary for the older subjects was  $a_1$ , the response time difference between the conditions would be  $c_1$  for young subjects and  $b_1$  for older subjects. Illustrating the same point in the bottom right panel of Figure 4, drift rate might be .4 for the easy condition and .3 for the hard condition for both young and older subjects, but the older subjects might be operating with boundary position a = .16, and the young with boundary position a = .12.

The right-hand panel of Figure 6 shows the distributions of response times that would result under this assumption—that is, the assumption that older subjects set wider boundaries. Again, the older subjects' response times are generally slower, and their distribution of response times is more spread out than that for the young subjects.

It should be mentioned here that changes in boundary positions have only modest effects on error rates if accuracy is reasonably high (e.g., 90% or better), effects that may not be detectable empirically. Looking for empirical differences in error rates to determine whether there are criterion shifts cannot be done in the absence of a model of processing, because, in some cases, large criterion shifts that produce a 100-msec or greater effect on response times may have little observable effect on accuracy, whereas in other cases there can be large (10%) effects on error rates (cf. Ratcliff & Rouder, 1998, Experiment 1).

Under both assumptions about how older subjects differ from young subjects, the distributions of response times shown in Figure 6 have about the same shape for older as for young subjects. Because the distributions have approximately the same shape, the resulting Brinley (Q–Q) plots will be approximately linear. Also, because the distribution is more spread out for the older subjects, the slope of the Brinley plot will be greater than 1. In other words, under both assumptions, the diffusion model makes predictions that are consistent with results from Brinley plot analyses.

The two illustrations we have given here are not the only possible differences between older and young subjects in drift rate and boundary positions. For example, a combination of drift and boundary position changes from young to older subjects would provide response time distributions similar to those in Figure 6. Also, the range of the nondecision component of response time  $(T_{\rm er})$  might be larger for older subjects than for young subjects, and the variability parameters might be larger for older subjects than for young subjects. In the framework of the diffusion model, there are multiple possible explanations for aging effects on response time, and which of these provides the best explanation can only be discovered by fitting the model to the data.

## SIMULATIONS USING THE DIFFUSION MODEL

The illustrations of how the effect of aging on response time might be viewed through the diffusion model have, up to this point, been illustrations of how the model works qualitatively. In this section, we show the results from simulations of the model.

Across the different simulations described below, we tested a number of plausible assumptions that might be made about how older subjects differ from young subjects in terms of the components of processing defined by the diffusion model. In each case, older subjects were assumed to differ from young subjects in their values of one or more of the parameters of the model. Variability was introduced into the parameter values to reflect variability across subjects. For example, the drift rate in an experimental condition might be chosen from a uniform distribution of drift rate values between, say, .2 and .6. For most of the simulations, the distributions of values of a



Figure 8. Simulated Brinley plot for older and young subjects with boundary positions, drift rates, and the nondecision component of response time varying across subjects. The parameter values are selected from uniform distributions and vary across the ranges shown in the figure.

parameter for older and young subjects were wide enough that they overlapped, with the consequence that fast older subjects could have faster response times than slow young subjects, just as in real data. Exact numerical solutions for predictions of response probability, mean response time, response time distributions, and so on are available for the diffusion model (see Ratcliff, 1978; Ratcliff et al., 1999). Given these solutions, exact values can be produced for the mean response time for a subject in an experimental condition, given the parameter values input to the model. Random variability arises from the random choices from the ranges of parameter values that represent the individual subjects, and it is these random choices that provide random variability in the simulations.

We describe the results of the simulations in terms of how the diffusion model produces linear Brinley plots with values of slope and intercept that match those that have been found in empirical data. For every simulation reported below, the slope was within 5% of the ratio of the standard deviation of older subjects' response times to the standard deviation of younger subjects' response times. The standard deviations were computed from the response times that were used to construct the Brinley plot, the response times for individual subjects in the first series of simulations, and the mean response times for experimental conditions in the second series of simulations.

## **A Single Experimental Condition**

We begin with simulations that model individual subjects' mean response times for just one experimental condition. First, it was assumed that older and young subjects set the same boundary positions (a = .16) but are differ-



Figure 9. Simulated Brinley plots for older versus young subjects, with differences among conditions modeled by consistent differences in drift rate (.05 difference between each condition). The top panel shows the results when older and young subjects differ only in boundary position, and the bottom panel shows the results when older and young subjects differ only in drift rate.

ent in their drift rates. That is, the experimental condition was more difficult for older subjects than for young subjects. The drift rate for each older subject was chosen randomly from a uniform distribution of drift rates between .2 and .4, and the drift rate for each young subject was chosen randomly from a uniform distribution of drift rates between .3 and .5. (Almost the same results would be obtained if the distributions of drift rates were normal.) The response time was then computed for each subject. The older subjects' response times were plotted against the young subjects' response times in the top panel of Figure 7, fastest older subject against fastest young subject, next fastest older subject against next fastest young subject, and so on (see Maylor & Rabbitt, 1994). The result shows the typical linear plot, with positive slope and negative intercept. The distributions of response times for both the older and the young subjects are symmetric, with a wider peak and shorter tails than a normal distribution, and the ratio of their standard deviations is within 5% of the slope of the Brinley plot.

The bottom panel of Figure 7 shows the Brinley plot when it is assumed that older subjects have the same drift rate as young subjects but set their response boundaries differently. For older subjects, the boundary positions were chosen from a uniform distribution with a range from .10 to .18, and for young subjects, the boundary positions were chosen from a uniform distribution with a range from .08 to .12. Again, the distributions of response times for both the older and the young subjects are symmetric, with a wider peak and shorter tails than a normal distribution.

Table 2 presents a more comprehensive picture of what happens to the slope and intercept of the Brinley plot as a function of the means and ranges of drift rates and boundary positions. In these simulations, only one parameter of the model was allowed to vary between subjects, either the boundary position or the drift rate. In each simulation, one mean response time was produced for each subject.

One simulation was conducted with 200 young subjects, with their drift rates varying between .3 and .7 and their other parameters as shown in the Table 2. To these young subjects, we compared seven different groups of older subjects (200 subjects per group), the groups varying in the mean and range of the distribution of their drift rates. The table shows the resulting intercepts and slopes of Brinley plots. The first three groups of older subjects each had the same range of drift values, but with increasingly lower means, and the last four groups had lower drift values and larger ranges. The resulting Brinley functions are linear, and the slope of the Brinley plot increases as the mean drift rate decreases, and as the slope increases, the intercept decreases.

A second simulation was conducted with 200 young subjects, with their boundary positions varying between .08 and .12 and the other parameters fixed as shown in the table. Again, response times for these young subjects were compared with response times for seven groups of older subjects, the older subjects' boundaries varying in the mean and range of their distributions. When the range of boundary positions is the same for older and young subjects, the Brinley slope is about 1. When the range is increased for older subjects, the slope increases. Because boundary position has a minimum value (zero), it is likely that boundary position becomes more variable as the boundaries move further apart. Therefore, we assumed that, as the mean value of the boundary position increased, the range increased. As is shown in the table, changes in boundary position produce linear Brinley functions with increasing slope when the mean boundary position and range increase; as the slope increases, the intercept decreases.

The results of these two sets of simulations show all the Brinley plot regularities. The Brinley plots are linear, with slopes and intercepts similar to those found em-



Figure 10. Simulated Brinley plots for older versus young subjects averaged across conditions, with differences among conditions modeled by differences in drift. The top panel shows the results from consistent differences in drift across conditions (the drift rate for each condition is .05 larger than the previous one). The bottom panel shows the result from adding a random value in the range from -.025 to .025 to the value of drift for each condition dition for each subject.

pirically. With either boundary position or drift rate varying, the results match empirical data. Thus, Brinley plots do not provide a test of which components of processing in the diffusion model would be responsible for the effects of aging.

In more realistic simulations, several parameters would vary simultaneously across individual subjects. Figure 8 shows the Brinley plots produced when older subjects are assumed to have smaller drift rates, wider boundaries, and larger baseline response times (larger values of  $T_{\rm er}$ ) than young subjects. We chose these three aspects of the model to vary because they are the ones that most often vary across individual young subjects (see, e.g., Ratcliff & Rouder, 1998; Ratcliff et al., 1999; the variability parameters, variability in drift across trials and variability in starting point, tend to affect error response times but have less effect on correct response times). For the sim-

ulations, for each individual subject, a value was chosen randomly for each of the three parameters from uniform distributions, using the ranges of values shown in the figure caption. The response time for the individual was generated from the model with these parameter values. The widths of the ranges of the values for all three of the parameters were assumed to be the same for older and young subjects, so that the only differences between older and young subjects were in the means of the parameters. The particular ranges chosen were extensions of the ranges of values found when the diffusion model is fit to empirical data for young subjects (Ratcliff, 1978; Ratcliff & Rouder, 1998; Ratcliff et al., 1999). Figure 8 shows four replications of the simulation. There are no systematic deviations from linearity. In each case, the function is linear, with a slope around 1.5, and the intercept is slightly negative or near zero. Thus, varying three of the parameters of the model to represent the effects of aging produces the same results as varying only drift rate or only boundary position. This reinforces the conclusion that Brinley plots cannot be used to identify what components of processing are affected by aging in the diffusion model.

## **Multiple Experimental Conditions**

The simulations just described show that the diffusion model produces Brinley plots that mimic experimental data when individual older subjects' response times are plotted against individual young subjects' response times, for a single condition of an experiment (i.e., a single value of difficulty). The diffusion model can successfully account for the empirical Brinley plot under any of several different assumptions and combinations of assumptions about which components of processing are affected by aging.

The more usual way to construct Brinley plots is to plot the mean response time for older subjects against the mean for young subjects, for each of a number of experimental conditions. To simulate these Brinley plots, we assumed that the experimental conditions differed in drift rate. This assumption corresponds to a within-subjects manipulation in which all the parameters of the model are fixed and only the quality of evidence coming from the stimulus changes across conditions. This means that subjects would be unable to change their speed-accuracy criteria (boundary positions in the diffusion model) as a function of stimulus condition. One simulation assumed that the only difference between older and young subjects was in boundary position settings, and the other assumed that the only difference was in drift rates (to parallel the simulations above).

For each simulation, there were 32 older and 32 young subjects. The top panel of Figure 9 shows the results when boundary positions were different for older and young subjects (with all other parameters set as shown in the figure). Older subjects' boundary positions were selected from a uniform distribution with a range of .08–.18, and young subjects' boundary positions were selected from a uniform distribution with a range of .06–.16. The drift



Figure 11. Average response time distributions for older and young subjects. The solid line is a transformation of the young subject response time distribution via the function  $2.59 \text{RT}_{v} + .600$ .

rate for the easiest experimental condition was selected from a uniform distribution with a range of .4–.6. The drift rates for the seven other conditions were successively .05 less than that for the largest drift rate. For each subject, a mean response time was computed for each condition. The response times for each condition were then averaged over subjects, and the condition means were used to construct the Brinley plot. The Brinley plot (Figure 9, top panel) is linear, with typical values of slope and intercept.

The bottom panel of Figure 9 shows the results with drift rate varying between young and older subjects. The drift rates for the easiest experimental condition were selected from uniform distributions with ranges of .4-.6 for older subjects and .6-.8 for younger subjects. The drift rates for the seven other conditions were successively .05 less than that for the largest drift rate. Again, the Brinley plot is linear, with typical slope and intercept.

Figure 10 shows the results of two simulations with drift rate, boundary position, and residual response time  $(T_{\rm er})$  all varying simultaneously between young and older subjects. The drift rate was assumed to change systematically across conditions, as above. Again, the slope and intercept (top panel of the figure) have typical values. For the simulation shown in the bottom panel, variability was added to the differences in drift rates for each condition. Instead of a consistent difference of .05 between the drift rates for each successive condition, a random number between -.025 and +.025 was added to the consistent drift value. The result (bottom panel of Figure 10) shows that the regular spacing of points in Figure 9

and Figure 10 (top panel) is a consequence of the regular spacing of drift rates; random variability makes that spacing less regular.

#### Summary

Across all the simulations, the diffusion model produced typical results for the effect of aging on response time: Brinley plots were linear, with positive slopes and negative intercepts, and the slopes were equal to the ratio of standard deviations for the older and young subjects' response times. The main conclusion is that aging differences can come from any of the several components of processing identified by the diffusion model or from combinations of the components. It would be easy to adjust the parameter values in any of the simulations to produce a specific value of slope—say, 1.5. In other words, the Brinley plot characteristics provide only weak constraints on modeling.

## OTHER EMPIRICAL RESULTS IN THE AGING AND RESPONSE TIME DOMAIN

Besides the Brinley plot regularities for mean response time for older versus young subjects, a number of other characteristics of response time data have received attention. In the following sections, we show that the diffusion model is consistent with these experimental results.

#### **Response Time Distributions**

G. A. Smith and Brewer (1995) found that average response time distributions (for correct responses) for older



Figure 12. Simulated quantile response times plotted against each other for changes in a single parameter in the model. (Note that the scales are different in the different panels, so that the slopes are not the same across panels.) These plots would correspond to plots of the quantiles in Figure 11 for older and young subjects against each other.

subjects slowed and spread, relative to those for young subjects (for a single experimental condition); that is, the older subjects' distribution could be produced from the young subjects' distribution by a transformation:  $RT_o = 1.12RT_y + 14$  msec. The average distributions over subjects were produced by Vincent averaging (see Ratcliff, 1979). The distributions had almost the same shape when scaled by the transformation and lay almost on top of each other.

We generated response time distributions with the diffusion model to check whether the distributions would fit the pattern G. A. Smith and Brewer (1995) found. The distributions G. A. Smith and Brewer examined had only relatively small differences in mean and spread of response times between young and older subjects. To more stringently test the diffusion model, we used parameter values that would give larger differences in mean and spread.

We generated distributions of response times, one distribution per subject, for 60 older subjects and 60 young subjects, using parameter values in the ranges shown in Figure 11, with older subjects having a larger boundary separation, a lower drift rate, and a larger nondecision component of response time. To average the response time distributions over subjects, quantile response times were computed for the response times for each subject, and these were averaged over subjects (just as in G. A. Smith & Brewer, 1995; see Ratcliff, 1979; Thomas & Ross, 1980). Figure 11 shows the average cumulative distributions for the young and older subjects. The solid line is the transformation (cf. G. A. Smith & Brewer, 1995) of the function for young subjects (see Figure 11). Just as in G. A. Smith and Brewer's study, the older subjects' distribution and the transformed young subjects' distribution lie on top of each other. In other words, the diffusion model can produce G. A. Smith and Brewer's empirical result.

We also examined whether variation in only one or another of the diffusion model parameters would give the correct predictions for the relative shapes of the young and older subjects' distributions. To show whether the distributions were (approximately) linear transformations of each other, as G. A. Smith and Brewer's (1995) data say they should be, we plotted the quantiles of the older and young response time distributions against each other. If the distributions are linear transformations of each other, the result should be a straight line. The leftmost panels show the results with drift rate alone varying between the older and the young (a drift rate of .4, in the top panel,



Figure 13. Standard deviations in response time as a function of drift rate and boundary position in the diffusion model.

plotted against a drift rate of .3, and a drift rate of .2, in the bottom panel, plotted against .3). The right-hand panels show the results with variation in boundary positions (boundary parameter a = .08 plotted against a = .12 in the top panel, and a = .16 plotted against a = .12 in the bottom panel). In both cases, the functions are linear. This means that the diffusion model can predict response time distributions that are linear transformations of each other.

Another regularity in the aging literature is that, when a particular quantile is plotted over experimental conditions, a straight line is obtained (see, e.g., Myerson et al., 1990; Smith, Poon, Hale, & Myerson, 1988). Because the quantiles shown in Figure 12 from the diffusion model are linear transformations of each other, plotting individual quantiles across a range of conditions (e.g., different drift rates) for older versus young subjects will produce linear functions.

## Standard Deviations in the Diffusion Model

The standard deviation in a single subject's response time for a single experimental condition has not been systematically studied but has occasionally been reported (e.g., Faust et al., 1999). In the diffusion model, the standard deviation can vary from changing as fast as mean response time changes across conditions to not changing at all as mean response time changes. Mean response time increases both when the distribution of response times shifts and when the distribution spreads. But the standard deviation increases only when the distribution spreads. If the distribution shifts only in position, not in spread, the standard deviation is constant. The behavior of the standard deviation in response time for the diffusion model is shown in Figure 13 as a function of boundary position and drift rate.

## Fast Subjects and Slow Subjects Versus the Group Mean

Hale and Jansen (1994; Balota & Ferraro, 1992; Zheng, Myerson, & Hale, 2000) plotted response times for fast young subjects and for slow young subjects against the mean response times for the whole group of young subjects. Specifically, they calculated the mean response time for the fastest half of the subjects in each condition of an experiment and plotted these means against the means for all the subjects in each condition, and they did the same thing for the mean response times for the slowest half of the subjects. The resulting two functions were both straight lines, with the line for the slow subjects' means having a larger slope and lower intercept than the line for the fast subjects' means. Thus, Brinley plot regularities can be found not just with differences between older and young subjects, but also for individual differences among young subjects. The Brinley plot slope was interpreted as an index of processing time.

We simulated Hale and Jansen's (1994) functions with the diffusion model, using 32 subjects and 16 experimental conditions. For each subject, the value for each of the three parameters of the model that were allowed to vary was selected randomly from a uniform distribution (v =.3-.55, a = .05-.15, and  $T_{er} = .30-.42$ ). To produce sys-



Figure 14. Simulated Brinley plots for fast and slow groups of subjects, compared with the mean for all subjects (cf. Hale & Jansen, 1994).

tematic differences among the 16 conditions, four values of drift rate (base plus .0, .04, .08, and .12) were crossed with four values of boundary position (base plus 0, .02, .04, and .06). In addition, variability in each parameter for each condition was added, the amount of variability chosen randomly from the range -.04 to .04 for drift and the range -.03 to .03 for boundary position. The overall mean for each subject was used to assign the subject to the fast or the slow group. Then, the mean response times for each condition for fast subjects and for slow subjects were plotted against group means, just as Hale and Jansen did. The results, shown in Figure 14, are straight lines, with a slope of 0.67 and an intercept of .068 for the fast group and a slope of 1.27 and an intercept of -.023 for the slow group (without variability in each condition in each parameter, the slope was 0.67 and the intercept .083 for the fast group, and the slope was 1.34 and the intercept -.060for the slow group). Thus, the diffusion model captures the same differences Hale and Jansen found with real data-that is, it captures individual differences within as well as across groups-and also accommodates hypotheses about Brinley plot slopes, such as the magnification hypothesis (Zheng et al., 2000).

## **Slopes Versus Intercepts of Brinley Plots**

Table 2 contains the slopes and intercepts of the Brinley plots for two sets of simulations of the diffusion model. Figure 15 graphs these slopes and intercepts against each other, as was done in Figure 3 for the data from Cerella (1985) and Faust et al. (1999). The functions in Figure 15

and Figure 3 are similar; they both have a negative slope and a positive intercept, with the slope being about 100 msec less than the intercept. This demonstrates the diffusion model's ability to capture the relationship between slopes and intercepts of the older/young subject data described by Cerella (1985) and predicted by the Q–Q plot analysis presented earlier.

## **GENERAL DISCUSSION**

Why is a Brinley plot typically linear, and why does it typically have a slope of about 1.5? In earlier research, the answer to this question has been that older subjects' response times are slower, relative to young subjects' response times, by a constant factor, typically about 1.5. In this article, we have used Q–Q theory to show that this answer is not correct. Instead, what Q–Q theory makes clear is that the plot is linear because the distributions of response times (across subjects or conditions) for older subjects' have about the same shape as the distributions for young subjects and that the slope is equal to the ratio of standard deviations of the older subjects' response times (across subjects or conditions) to the young subjects' response times.

This reinterpretation makes the Brinley regularities easier to understand. The range of older subjects' response times is wider than the range of young subjects' response times. If we believe that older subjects are not too different in their processing from young subjects, the older subjects' range being only about 1.5 times wider is



Figure 15. Intercepts plotted against slopes from the diffusion model simulations shown in Table 2.

not too surprising. For example, if young subjects' response times ranged from 500 to 800 msec, the Brinley slope would be 1 if older subjects' response times also ranged from 500 to 800 msec (assuming that the distributions had the same shape). Or, if younger subjects' times ranged from 500 to 800 msec and older subjects' times ranged from 500 to 1,400 msec, the Brinley slope would be 3. Thus, instead of asking whether a model can produce linear Brinley functions, it is only necessary to ask whether the distribution of mean response times across older subjects has the same shape as the distribution across young subjects, and whether and by how much the spread of the older subjects' distribution is wider. The interpretation in terms of distributions provides a simpler meeting point between models of processing and aging effects. The connection between standard deviations in response times and a model is more easily understood than the connection between the Brinley function and a model.

Under the Q–Q interpretation of Brinley plots, the target of modeling is the relative shapes and spreads of the distribution of mean response time across conditions or subjects of older versus young groups, not the fact that the slope typically has a value of about 1.5. In moving attention away from the slope value as the theoretically most relevant piece of data, the Q–Q analysis agrees with the general conclusion from a recent debate on the utility of Brinley plots (Cerella, 1994; Fisk & Fisher, 1994; Myerson et al., 1994; Perfect, 1994). The conclusion was that Brinley plots provide only weak constraints either on models of processing or on models of the effects of aging. It was argued that any of a number of assumptions about changes as a function of aging are probably capable of explaining the Brinley patterns of results.

The fact that the Brinley plot, whether considered, as it was in the past, as a measure of general slowing or considered via O-O analysis as a measure of relative distributions, does not constrain models is illustrated by the diffusion model simulations we carried out. For the diffusion model, wider distributions of response times for older than for young subjects can come about because drift rates have a wider range for older subjects, because older subjects set more conservative boundary positions, or because of some combination of these factors, as well as because of changes in the nondecisional component of processing. Moreover, there is no single set of values for any of these parameters that is required to get a particular Brinley slope and intercept; many sets of parameters values can fit any particular Brinley plot. It is almost trivial for the diffusion model to simulate data that have all the usual Brinlev regularities.

The critical tests among models will involve aspects of the data other than the Brinley regularities. For the diffusion model, the critical tests involve the joint behavior of response time and accuracy, the way the shape of the response time distributions changes as a function of drift rate or boundary position, and the relative speeds of correct and error responses (see Ratcliff & Rouder, 1998, 2000; Ratcliff et al., 1999). Because the model is forced to fit all these measures simultaneously, the model is severely constrained. Jointly fitting all these dependent variables can cause the model to fail. But if it succeeds in fitting the data adequately, the behavior of the parameters of the model can be used to interpret the effects of aging.

This approach offers the potential of being able to discriminate older from young subjects in multiple ways. One question is whether the quality of the information that drives the decision process (drift rate) is different or more variable across trials for older subjects than for young subjects. Another question is whether older subjects set their decision criteria (speed–accuracy criteria) more conservatively than do young subjects, and how variable are the criteria across trials. Another possibility is that the nondecisional component of response time is longer for older subjects. Increases in task difficulty might affect older subjects more than young subjects, and this might show up in any or all of the components of processing just listed.

The agenda set by this type of processing model is to examine performance of older versus young subjects across a range of different kinds of cognitive tasks. The interesting questions are what components of processing are affected differentially for the older subjects by the changes in task and what components hold constant. For example, if an individual older subject is more conservative in criteria setting than the average older subject in one task, is he/she more conservative in other tasks? If the quality of information extracted from the stimulus (drift rate) is lower in one task, is it also lower in other tasks, and in what kinds of other tasks? This approach will not only be able to identify differences between older and young groups but will also be able to examine consistency of components of processing within and across individuals within a group.

In sum, it is our hope that the theoretically based approach to the study of aging offered by the diffusion model and other such models will lead to a qualitatively deeper understanding of the effects of aging on response time. The earlier, more descriptive models have pointed out regularities in the data, and we think the stage is set for theory-based examinations of the effects of aging on response time.

#### REFERENCES

- BALOTA, D. A., & FERRARO, F. R. (1992, November). What is unique about age in age-related general slowing? Paper presented at the 33rd Annual Meeting of the Psychonomic Society, St. Louis.
- BASOWITZ, H., & KORCHIN, S. J. (1957). Age differences in the perception of closure. Journal of Abnormal & Social Psychology, 54, 93-97.
- BRINLEY, J. F. (1965). Cognitive sets, speed and accuracy of performance in the elderly. In A. T. Welford & J. E. Birren (Eds.), *Behavior, aging and the nervous system* (pp. 114-149). Springfield, IL: Thomas.
- CERELLA, J. (1985). Information processing rates in the elderly. Psychological Bulletin, 98, 67-83.
- CERELLA, J. (1990). Aging and information-processing rate. In J. E. Birren & K. W. Schaie (Eds.), *Handbook of the psychology of aging* (3rd ed., pp. 201-221). San Diego: Academic Press.
- CERELLA, J. (1991). Age effects may be global, not local: Comment on Fisk and Rogers (1991). Journal of Experimental Psychology: General, 120, 215-223.
- CERELLA, J. (1994). Generalized slowing in Brinley plots. Journal of Gerontology: Psychological Sciences, 49, 65-71.
- CHAMBERS, J. M., CLEVELAND, W. S., KLEINER, B., & TUKEY, P. A. (1983). Graphical methods for data analysis. Boston: Duxbury.

- DOSHER, B. A. (1984). Discriminating preexperimental (semantic) from learned (episodic) associations: A speed–accuracy study. *Cognitive Psychology*, **16**, 519-555.
- DRAPER, N. R., & SMITH, H. (1981). *Applied regression analysis*. New York: Wiley.
- FAUST, M. E., BALOTA, D. A., SPIELER, D. H., & FERRARO, F. R. (1999). Cognitive slowing, cognitive speed, and response latency. *Psychological Bulletin*, **125**, 777-799.
- FISHER, D. L., & GLASER, R. A. (1996). Molar and latent models of cognitive slowing: Implications for aging, dementia, depression, development, and intelligence. *Psychonomic Bulletin & Review*, 3, 458-480.
- FISK, A. D., & FISHER, D. L (1994). Brinley plots and theories of aging: The explicit, muddled, and implicit debates. *Journal of Gerontology: Psychological Sciences*, **49**, 81-89.
- HALE, S., & JANSEN, J. (1994). Global processing-time coefficients characterize individual and group differences in cognitive speed. *Psychological Science*, 5, 384-389.
- LAMING, D. R. J. (1968). Information theory of choice reaction time. New York: Wiley.
- MAYLOR, E. A., & RABBITT, P. M. A. (1994). Applying Brinley plots to individuals: Effects of aging on performance distributions in two speeded tasks. *Psychology & Aging*, 9, 224-230.
- MYERSON, J., HALE, S., WAGSTAFF, D., POON, L. W., & SMITH, G. A. (1990). The information-loss model: A mathematical theory of agerelated cognitive slowing. *Psychological Review*, **97**, 475-487.
- MYERSON, J., WAGSTAFF, D., & HALE, S. (1994). Brinley plots, explained variance, and the analysis of age differences in response latencies. *Journal of Gerontology: Psychological Sciences*, **49**, 72-80.
- NEBES, R. D., & MADDEN, D. J. (1988). Different patterns of cognitive slowing produced by Alzheimer's disease and normal aging. *Psychology & Aging*, 3, 102-104.
- PACHELLA, R. G. (1974). The interpretation of reaction time in information processing research. In B. Kantowitz (Ed.), *Human information processing: Tutorials in performance and cognition* (pp. 41-82). New York: Halstead.
- PERFECT, T. J. (1994). What can Brinley plots tell us about cognitive aging? Journal of Gerontology: Psychological Sciences, 49, 60-64.
- 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. (1981). A theory of order relations in perceptual matching. *Psychological Review*, 88, 552-572.
- RATCLIFF, R. (1985). Theoretical interpretations of speed and accuracy of positive and negative responses. *Psychological Review*, 92, 212-225.
- RATCLIFF, R. (1988). Continuous versus discrete information processing: Modeling the accumulation of partial information. *Psychological Review*, 95, 238-255.
- RATCLIFF, R., & MURDOCK, B. B., JR. (1976). Retrieval processes in recognition memory. *Psychological Review*, 83, 190-214.
- RATCLIFF, R., & ROUDER, J. F. (1998). Modeling response times for two-choice decisions. *Psychological Science*, 9, 347-356.
- RATCLIFF, R., & ROUDER, J. F. (2000). A diffusion model account of masking in letter identification. *Journal of Experimental Psychol*ogy: Human Perception & Performance, 26, 127-140.
- RATCLIFF, R., VAN ZANDT, T., & MCKOON, G. (1999). Connectionist and diffusion models of reaction time. *Psychological Review*, **106**, 261-300.
- REED, A. V. (1973). Speed–accuracy trade-off in recognition memory. *Science*, 181, 574-576.
- REED, A. V. (1976). List length and the time course of recognition in immediate memory. *Memory & Cognition*, **4**, 16-30.
- SALTHOUSE, T. A., & SOMBERG, B. L. (1982). Isolating the age deficit in speeded performance. *Journal of Gerontology*, **37**, 59-63.
- SILVERMAN, I. (1963). Age and the tendency to withhold responses. *Journal of Gerontology*, **18**, 372-375.
- SMITH, G. A., & BREWER, N. (1995). Slowness and age: Speed–accuracy mechanisms. *Psychology & Aging*, 10, 238-247.
- SMITH, G. A., POON, L. W., HALE, S., & MYERSON, J. (1988). A regular relationship between old and young adults' latencies on their best, average and worst trials. *Australian Journal of Psychology*, **40**, 195-210.
- SMITH, P. L., & VICKERS, D. (1988). The accumulator model of two-choice discrimination. Journal of Mathematical Psychology, 32, 135-168.

#### 24 RATCLIFF, SPIELER, AND MCKOON

- 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 & Performance*, 22, 461-479.
- SPIELER, D. H., BALOTA, D. A., & FAUST, M. E. (in press). Levels of selective attention revealed through analyses of response time distributions. *Journal of Experimental Psychology: Human Perception & Performance.*
- STRAYER, D. L., & KRAMER, A. F. (1994). Aging and skill acquisition: Learning-performance distinctions. *Psychology & Aging*, 9, 589-605.
- THOMAS, E. A. C., & ROSS, B. H. (1980). On appropriate procedures for combining probability distributions within the same family. *Journal* of Mathematical Psychology, **21**, 136-152.
- THORNDIKE, E. L., BREGMAN, E. O., TILTON, J. W., & WOODYARD, E. (1928). *Adult learning*. New York: Macmillan.
- WICKELGREN, W. A. (1977). Speed–accuracy tradeoff and information processing dynamics. Acta Psychologica, 41, 67-85.
- ZHENG, Y., MYERSON, J., & HALE, S. (2000). Age and individual differences in visuospatial processing speed: Testing the magnification hypothesis. *Psychonomic Bulletin & Review*, 7, 113-120.

## APPENDIX

## Linear Regression With Variability in Both the x and the y Values

An issue that should be considered is the extent to which variability in both the x and the y values in linear regression applied to the Brinley plot affects the estimate of the slope. In a typical Brinley plot, we would expect the ordering of mean response times across experimental conditions to be the same for young and for older adults. To test what happens when variability is added to the x and y values, we assumed that the orders were the same for older and for young subjects and that the underlying distributions were the same (uniform) and then added random variability to both the x and the y values.

Table A1 shows the assumed ranges of response times and the results of adding variability. In the top half of the table, the older subjects' mean response times for 10 experimental conditions range from 475 to 1,150 msec, with 475 msec the mean for the fastest condition, 1,150 msec the mean for the slowest condition, and the other means located at equally spaced steps of 50 msec. The young subjects' means range from 450 to 900 msec in steps of 75 msec. To each of the means, noise was added in the form of a a random number chosen from a normal distribution, with the mean of zero and the standard deviation shown in the third column of the table. The resulting ratios of standard deviations and Brinley slopes are shown in the fourth and sixth columns. As long as the added variability is less than 50 msec, the ratio and slope are not affected, to within standard error estimates. This corresponds to a correlation between x and y values  $(r^2)$  of .9 or greater. When the amount of noise is very large, the condition means become uncorrelated, the Brinley slope goes to zero, and the ratio of standard deviations becomes 1. In real data, the correlation between the condition means is usually above .9, so the effect of variability in the x and y values does not affect the estimate of slope significantly. The slope was also not affected when the Brinley slope was set to 2.0, shown in the bottom half of Table A1.

The results in Table A1 show an interesting pattern of bias. In the top half of the table, the true ratio of standard deviations is 1.5, but the Brinley slope estimated from linear regression falls to around 1.1 as 100 msec of noise is added to the true values. The ratio of standard deviations is less biased, but still falls to around 1.4. The bottom half of the table, with a true slope of 2.0, shows the same effect, the Brinley slope falling to 1.44 and the ratio of standard deviations falling to 1.85.

This bias away from the true value is a topic of some research in statistics. The problem is that the model for linear regression is  $y = a + bx + \varepsilon_i$  (which is the standard expression) and variability is assumed to only occur in the y value. But, when there is also variability in the x value,  $x = X + \xi_i$ , where X is the true x value, the estimate of slope from the linear regression is biased. A number of different approaches to this problem have been proposed, and which approach is best depends on whether the two sources of variability are independent, correlated, constant, or vary as a function of their values, and so on. For example, Draper and Smith (1981, p. 124) list conditions under which the usual linear regression is adequate and present a method for other situations (we tried that method, and it was not superior to linear regression in our examples from Table A1).

The conclusion is that when there is a moderate amount of variability in both the x and the y values (corresponding to a correlation coefficient less than .9), the slope estimated from linear regression is likely to be biased away from the true slope. Thus, when any hypotheses are tested about equality of slopes across conditions, it is possible to obtain differences in the estimates when, in fact, there are no true differences in variability in the conditions. One solution that reduces the bias (but does not eliminate it) is to use the ratio of standard deviations in the y and x scores to estimate the slope when the distributions appear to come from approximately the same family.

of the Brinley Plot, and Correlation Between Young and Old Reaction Times (RTS)										
Range of Old RTs	Range of Young RTs	<i>SD</i> of Noise Added	Ratio of <i>SD</i> s of Old to Young RTs	SE in Ratio of SDs	Slope of Brinley Plot	<i>SE</i> in Slope of Brinley Plot	Correlation	SE in Correlation		
475-1,150	450-900	0	1.50	0	1.50	0	1.00	0		
475-1,150	450-900	25	1.52	0.10	1.50	0.10	.97	.02		
475-1,150	450-900	50	1.47	0.19	1.38	0.19	.89	.05		
475-1,150	450-900	75	1.48	0.26	1.27	0.24	.75	.11		
475-1,150	450-900	100	1.39	0.29	1.09	0.27	.63	.17		
500-1,400	450-900	0	2.00	0	2.00	0	1.00	0		
500-1,400	450-900	25	1.98	0.12	1.94	0.12	.97	.01		
500-1,400	450-900	50	1.96	0.23	1.84	0.22	.88	.06		
500-1,400	450-900	75	1.91	0.27	1.68	0.22	.77	.10		
500-1,400	450-900	100	1.85	0.42	1.44	0.39	.63	.16		

Table A1 Simulations for Values of the Ratio of Standard Deviations in Scores, Slope of the Brinley Plot, and Correlation Between Young and Old Reaction Times (RTs)

Note—Ten steps between the limits in the ranges, linear steps. The standard deviation in the young RTs is 151, and the standard deviation in the old RTs is 227 for the 475–1,150 range and 303 for the 500–1,400 range. The noise added is a normally distributed random variable with standard deviation in column 3.

(Manuscript received May 21, 1998; revision accepted for publication April 21, 1999.)