

RTSYS: A DOS application for the analysis of reaction time data

ANDREW HEATHCOTE

University of Newcastle, Newcastle, New South Wales, Australia

RTSYS is a menu-driven DOS application for the manipulation, analysis, and graphical display of reaction time data. It can be used either in a single-task environment under DOS, with access to a set of operating system commands, or as an application under Windows. All functions have context-sensitive help. RTSYS fits the ex-Gaussian distribution to reaction time data without the difficulties usually associated with numerical parameter estimation. Distribution fitting and flexible censoring and rescaling options allow RTSYS to address the problems of reaction time distribution skew and outlying responses with reasonable sample sizes. RTSYS can automatically process multiple input files from experiments with arbitrary designs and produce formatted output of statistics for further processing by graphical and inferential statistical packages. The present article reviews and explains techniques used by RTSYS and provides an overview of the operation of the program.

RTSYS is a DOS application, written in Turbo Pascal 6.0, that calculates distribution statistics for reaction time (RT) data. It also provides facilities for censoring and rescaling data. Statistics calculated include the number of RTs, percent of RTs censored, the median, mean, variance, and a nonparametric measure of skew [i.e., (mean - median)/standard deviation]. Because RT data are often scored as correct or incorrect, where incorrect RTs are relatively infrequent, percent error and mean error RT are also calculated. RTSYS fits the ex-Gaussian distribution and reports distribution parameters μ , σ , and τ . Fitting the ex-Gaussian also produces chi-squared and likelihood goodness-of-fit statistics. Goodness of fit can be inspected visually by plotting RT histograms and superimposed fitted ex-Gaussian distributions.

RTSYS will typically be used to convert raw data files collected from experiments into files of statistical parameters that are used as input to inferential and graphical applications. One of RTSYS's most useful features is its ability to automatically process arbitrary factorial and nonfactorial between- and within-subject designs. Design cells can be collapsed both by mixing data from different factor levels or by Vincent averaging. Vincent averaging (usually performed over a subjects factor) approximately preserves distribution shape and is useful for obtaining ex-Gaussian fits where design cells contain too few observations for sufficiently precise parameter estimates. Vin-

cent probability histograms allow visual inspection of the shape of the averaged distribution and the goodness of fit of a fitted ex-Gaussian distribution.

Once you have specified a within-subject design and the names of the data files corresponding to each subject, RTSYS automatically calculates statistics for each subject and design cell. Individual parameters for each subject and subject averages for each cell can be viewed on screen, printed, or output to a parameter file. Parameter file format, including delimiter and missing value characters, and the order of design cells can be controlled to conform to the requirements of graphical and inferential applications.

RTSYS is accompanied by a manual that provides an extensive overview of the program and examples of how to use options and organize analyses. Example data files are provided so that the user can practice using RTSYS. RTSYS also has on-line context-sensitive help for all options as well as an introductory help that gives an overview of the application. The first section of this article reviews and explains the analysis techniques used by RTSYS. The remainder of the article describes the user interface and hardware requirements of RTSYS.

WHY SHOULD YOU USE RTSYS?

Two phenomena make the analysis of RT data problematic. First, RT data can contain fast and slow outlying values caused by anticipation and distraction. If the proportion of outliers is sufficient, interpretation of results may be confounded. Second, RT distributions are usually positively skewed (e.g., Figure 1; but see Luce, 1986, p. 117, for a counterexample with simple RT to an intense auditory stimulus). Skew creates problems of interpretation for descriptive statistics. For instance, an independent variable may affect the mean and median differently by changing the degree of skew. Significantly skewed or

The development of RTSYS began during the author's doctoral research, which was supported by a Canadian Commonwealth Fellowship. Thanks to Doug Mewhort and Richard Heath for comments on the first draft of this manuscript, to Trish Van Zandt, Steve Link, and an anonymous reviewer for their comments, and to the numerous people who have tested versions of RTSYS. Thanks also to Sarah Ogilvie for proof-reading the manuscript and providing a quiet place in Melbourne to finish this project. Correspondence should be addressed to A. Heathcote, Psychology Department, University of Newcastle, University Drive, Newcastle, NSW 2308, Australia.

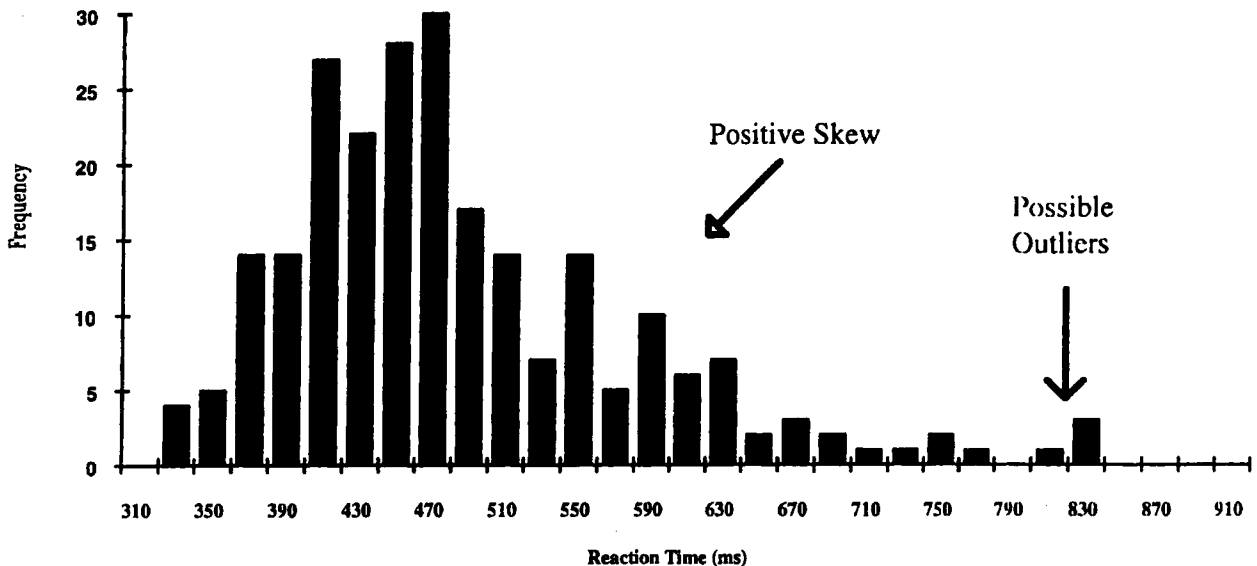


Figure 1. A histogram of RT data exhibiting a skewed distribution and potential outliers.

asymmetric data also violate an assumption made by most parametric tests—that variability in data is normal and hence distributed symmetrically.

Skew is commonly ignored, or it is removed by nonlinearly transforming the measurement scale, for instance, from milliseconds to $\log(\text{millisecond})$ or millisecond^{-1} (e.g., Box & Cox, 1964, 1982). While this strategy is useful to ensure that data conform to a normal distribution assumption, Ratcliff and Murdock (1976), Hockley (1984), Mewhort, Braun, and Heathcote (1992), and Heathcote, Popiel, and Mewhort (1991) all demonstrate that the magnitude of skew can contain information about the effect of experimental manipulations. Hence, transformation may lead inferential tests to miss potentially important effects. Furthermore, analysis of untransformed RT takes advantage of a natural and interpretable ratio scale, time. Nonlinear transformations of time produce scales that are not so easily interpretable, with the exception of the inverse transform that results in a speed scale.

Solutions to the problem of outliers usually rely on removing or censoring observations. Criteria for censoring are, however, problematic because real data are almost inevitably rejected along with spurious data. The amount of real data lost as a function of the exclusion criterion is usually unknown, so one must choose criteria without being able to weigh the cost of losing real data against the benefit of excluding spurious data. Despite these difficulties, outliers cannot be ignored, especially in paradigms where frequent distraction or anticipation is known to occur. Beginning with Tukey's (1960) seminal paper on the contaminated normal model, there have been numerous demonstrations of the potentially large influence of even a small percentage of outliers on statistical parameter estimates.

A further practical problem, generated in part by attempted solutions to the difficulties outlined above, is the plethora of statistics that must be calculated and interpreted by the conscientious RT researcher. Location (e.g., mean and median), scale (e.g., variance), and asymmetry (e.g., skew) measures, and potentially a variety of nonlinearly rescaled versions, need to be calculated. Percent error is useful to address the problem of speed-accuracy tradeoff (Johnson, 1939; Link, 1982; Ollman, 1966; Pachella, 1974; Yellot, 1967). Mean error RT is useful to determine if errors are due to anticipation or distraction. To compound the problem, each censoring scheme doubles the number of statistics to be analyzed.

While RTSYS does not provide a panacea, it encourages researchers to explore possible solutions to the problems outlined with a range of techniques. Through quantifying RT distribution shape, RTSYS can reveal structure within RT data not evidenced by conventional analyses. The following sections provide relatively non-technical explanations and illustrations of the techniques employed by RTSYS. The first section deals with the issue of RT distribution skew. The second section deals with Vincentizing, a useful technique when the number of observations is low. The third section deals with the issue of outliers.

Estimating RT Distribution Skew

The principal difficulty in determining skew is finding an *efficient estimator*. Efficiency is related to the variation of estimates across samples. When efficiency is low, estimates vary widely, especially with small samples. Prior to the widespread availability of computers, the ease of computation of an estimator was also a consideration. While the computational burdens of the techniques

outlined are quite heavy, RTSYS has been optimized to allow adequate performance even on an XT or AT.

Kendall and Buckland's (1967) *Dictionary of Statistical Terms* defines skewness as "an older and less preferable term for asymmetry, in relation to a frequency distribution." Most statistical packages describe skewness with a statistic called the *third shape factor* or *moment ratio* (often simply called *skew*): $\alpha_3 = (\mu_3/\mu_2)^{2/3}$. The variance, μ_2 , and the third central moment, μ_3 , are estimated by:

$$\hat{\mu}_2 = \frac{\sum_{i=1}^n (x_i - \bar{x})^2}{n - 1}$$

and

$$\hat{\mu}_3 = \frac{\sum_{i=1}^n (x_i - \bar{x})^3}{n - 1}.$$

The $2/3$ exponent and μ_2 components of α_3 are scale factors. It is μ_3 that measures the asymmetry of the distribution. As can be seen from its estimator, μ_3 is negative when the distribution is left skewed (mean < median for unimodal distributions) and positive when the distribution is right skewed (mean > median, the usual case with RT data).

The formulas used to estimate α_3 are justified using the method of moments. The method of moments does not require specification of a particular theoretical distribution, although it does require that moments of the desired order exist (cf. the Cauchy distribution for which moments are undefined), and provides formulas that are simple to compute. However, Ratcliff (1979) criticized estimation of skew in RT data by the method of moments on the grounds that it is inefficient (i.e., estimates are very variable, requiring tens of thousands of observations to become sufficiently precise) and not robust (i.e., it is sensitive to outliers). RTSYS calculates the standard deviation, using the method-of-moments formula, but it does not calculate the third central moment or skewness using the method-of-moments formula because of its low efficiency.

Ratcliff suggested that, instead of the method of moments, maximum likelihood estimation be used to determine skewness. Maximum likelihood estimation is more efficient than the method of moments. In fact, it is consistent and efficient, and it has the highest asymptotic (large sample) efficiency of any unbiased estimation method (see Cox & Hinkley, 1974, chap. 9). Its drawback is that a particular theoretical form of the RT distribution must be assumed. If the assumption is wrong, estimates will be biased. It is also computationally intensive, requiring optimization or search rather than providing a simple formula such as that for μ_3 .

The Ex-Gaussian Distribution

The theoretical distribution used by Ratcliff (1978, 1979; Ratcliff & Murdock, 1976) as a model of RT distribution is the sum of independent Gaussian (normal) and exponential random variables, named the ex-Gaussian distribution by Burbeck and Luce (1982). The ex-Gaussian distribution has three parameters: μ , the mean of the normal component; σ , the standard deviation of the normal component; and τ , the mean of the exponential component. The parameters of the components are related in a straightforward manner to the central moments of the overall distribution, the mean (μ_1), the variance (μ_2), and the third central moment (μ_3):

$$\mu_1 = \mu + \tau \quad (1)$$

$$\mu_2 = \sigma^2 + \tau^2 \quad (2)$$

$$\mu_3 = 2\tau^3 \quad (3)$$

Hence, the overall mean is determined by the location of both normal and exponential components, the overall variance is determined by both the variance of the normal and exponential components (the variance of an exponential is τ^2), and μ_3 is determined solely by the exponential component, because the normal distribution is symmetrical and hence makes no contribution to the asymmetry of the ex-Gaussian.

Figure 2 illustrates an ex-Gaussian distribution and its Gaussian and exponential components. The curves are called *probability density functions* (pdfs) because the probability of an RT in a given interval is the area




Figure 2. Probability density functions for (a) a normal distribution with $\mu = 500$, $\sigma = 50$, (b) an exponential distribution with $\tau = 100$, and (c) the resulting ex-Gaussian distribution. Parameter values are typical of those observed by Heathcote, Popiel, and Mewhort (1991). Units are expressed in milliseconds.

under the curve in that interval. For example, the probability of a 500-msec RT measured on a system with millisecond resolution (i.e., the response occurred in the interval between the 499th and 500th system clock tick) is the integral (i.e., area) of the pdf from 499 to 500 msec. The pdfs for the exponential and Gaussian components are given by the following equations:

$$pd_{\text{Ex}}(\text{RT} | \tau) = \frac{1}{\tau} e^{-\frac{1}{\tau}\text{RT}}$$

and

$$pd_{\text{G}}(\text{RT} | \mu, \sigma) = \frac{1}{\sigma\sqrt{2\Pi}} e^{-\frac{1}{2}\left(\frac{\text{RT}-\mu}{\sigma}\right)^2}$$

The pdf for the ex-Gaussian is calculated by taking the convolution of the above equations.

Convolution of continuous random variables is most easily explained using the analogous operation for discrete random variables. A discrete random variable is described by a probability function (pf), the discrete analogue of the continuous pdf. The pf gives the probability of each of a set of discrete outcomes. For a fair coin, for example, the pf is $p(x) = .5$, where $x = 0$ for a tail (T) and $x = 1$ for a head (H). Now consider the probability function for the sum of two coin tosses, $p_2(x)$, where $x = 0$ for (T,T), $x = 1$ for (T,H) or (H,T), and $x = 2$ for (H,H). The summation of coin-toss outcomes is analogous to the summation of outcomes for the continuous Gaussian and exponential components of the ex-Gaussian. Given that the coin tosses are independent, the probability for any pair of outcomes is $1/4$, the product of the pf value for each outcome. To finish constructing the pf, we need only add the probabilities of pairs with identical outcomes—in this case, the single outcome $x = 1$. Hence, $p_2(0) = 1/4$, $p_2(1) = 1/2$, and $p_2(2) = 1/4$.

Convolution in the continuous case involves analogous operations of multiplying pdfs and summing densities for identical outcomes. Summation of identical outcomes is replaced by integration over a variable z , where the outcome of one pdf is set to z and the other pdf to $(\text{RT} - z)$, so that the sum of outcomes is $z + (\text{RT} - z) = \text{RT}$. In the following equation, the integration is performed over values where the exponential distribution has nonzero probability density $(0, \infty)$.

$$\begin{aligned} pd_{\text{Ex-G}}(\text{RT} | \mu, \sigma, \tau) &= \int_0^{\infty} pd_{\text{Ex}}(z, \tau) pd_{\text{G}}[(\text{RT} - z), \mu, \sigma] dz \\ &= \frac{\sigma^2}{\tau\sqrt{2\Pi}} e^{-\frac{\text{RT}-\mu}{\tau}} \int_{-\infty}^{\left(\frac{\text{RT}-\mu}{\sigma} - \frac{\sigma}{\tau}\right)} e^{-\frac{y^2}{2}} dy. \end{aligned} \quad (4)$$

Unfortunately, the convolution integral does not have a closed form solution. However, the result can be ex-

pressed as above using the standard normal integral, which has a number of accurate and computationally cheap numerical approximations.

The ex-Gaussian was first proposed by McGill (1963) as a model of RT data where the right tail of the RT distribution seemed to be exponential. McGill suggested that the exponential distribution reflects the residual component of RT (the sum motor and perceptual and other nondecision stages), with the Gaussian distribution due to decision latency. However, Luce (1986) points out that a review of physiological mechanisms by Meijers and Eijkman (1974) found that the motor component of the residual contributes very little variability. Since the exponential component is typically the main contributor of variability to the ex-Gaussian, Luce suggests that it should not be identified with the low variance residual component of RT.

The opposite attribution of components was suggested by Hohle (1965). He used the central limit theorem to argue that residual time, because it is the sum of a number of component processes, is Gaussian. However, as the sum of Gaussians is Gaussian, part of the Gaussian component could be due to decision processes even if Hohle is correct. Later work suggests that both the Gaussian and exponential components of fitted ex-Gaussian distributions can be functions of experimental manipulations that should influence the decision stage (e.g., Heathcote et al., 1991; Hockley, 1984), rejecting a unique identification of the components of the ex-Gaussian with residual and decision times.

Luce (1986) notes that the fit of the ex-Gaussian to RT data is "surprisingly accurate" (p. 100). For example, Ratcliff and Murdock (1976) found that the ex-Gaussian provided a better fit to data from recognition memory experiments than did gamma and lognormal distributions. Heathcote (1995a) argues that the ex-Gaussian is a good model of RT data because the Gaussian component can absorb additive Gaussian measurement error. Hence, the ex-Gaussian can model variability associated with the measurement process as well as variability intrinsic to the cognitive mechanisms producing the observed RT.

Estimation of skewness through maximum likelihood fitting of the ex-Gaussian was implemented in RTSYS because of the ex-Gaussian's empirically demonstrated ability to accurately model RT distribution. However, it is hoped that, in the future, fitting of alternative distributions will be implemented. Luce (1986) provides a comprehensive review of theoretically motivated RT distribution models. Candidates for implementation in RTSYS include the displaced gamma, lognormal, inverse Gaussian, and Weibull distributions, all of which have a simple three-parameter form and have received attention in the psychological literature. Fitting of these distributions through maximizing likelihood is not always as easy as for the ex-Gaussian because of nonregular parameters. Regular parameters are asymptotically normally distributed. When a parameter is not regular, it has poor estimation properties. Hence, another reason to

prefer the ex-Gaussian is that numerical estimation through likelihood maximization is stable and robust.

Estimating the Ex-Gaussian Parameters

This section describes estimation of the ex-Gaussian parameters by likelihood maximization and compares it with estimation using the method of moments. The likelihood of an observation is identical to its probability density for a continuous distribution, or probability for a discrete distribution. However, while identical mathematically, the emphasis is shifted from the probability density of a particular observation given a set of distribution parameters, $pd(RT|\mu, \sigma, \tau)$, to the probability density of a particular set of distribution parameters given the observation, $L(\mu, \sigma, \tau|RT)$. More commonly, likelihood is considered as the joint probability density of a set of observations, $RT = (RT_1, RT_2, \dots, RT_n)$. Given that the observations are independent, the likelihood is given by the product of the probability densities of the individual observations.

$$L(\mu, \sigma, \tau|RT) = pd(RT_1|\mu, \sigma, \tau) pd(RT_2|\mu, \sigma, \tau) \dots pd(RT_n|\mu, \sigma, \tau).$$

For fixed RT , the likelihood gives the probability density of particular values of the parameters. Maximum likelihood estimation selects the parameter values that have the greatest probability density—that is, the parameter values that are most likely given the data set.

In some cases, maximum likelihood estimation leads to closed form solutions for the parameters as a function of the data. For example, least squares estimation for linear models with additive, constant variance Gaussian error is equivalent to the maximum likelihood estimate. For the ex-Gaussian distribution, however, closed form solutions are not available, so maximum likelihood estimation is carried out by search. Search involves evaluating the likelihood for different values of the parameters until the best values are found.

Because the probability density values for each observation are usually small, their product, the likelihood, can become so small that it is difficult to represent in floating point form on a computer. To avoid representational difficulties, search is usually carried out using log likelihood $[l(RT|\mu, \sigma, \tau)]$ rather than likelihood itself. Logarithms are monotonic, and, therefore, order-preserving, functions. Hence, the parameter values obtained by maximizing likelihood are the same as the parameter values for the maximum log likelihood. Logarithms also help to simplify the expression for the likelihood, because they

convert products to sums. The ex-Gaussian log likelihood is given by:

$$l_{Ex-G}(RT|\mu, \sigma, \tau) = \ln L_{Ex-G}(RT|\mu, \sigma, \tau) = -n \left(\ln \sqrt{2\pi} + \ln \tau - \frac{\sigma^2}{2\tau^2} - \frac{\mu}{\tau} \right) - \sum_{i=1}^n \left[\frac{RT_i}{\tau} - \Phi \left(\frac{RT_i - \mu}{\sigma} - \frac{\sigma}{\tau} \right) \right].$$

$\Phi(x)$ is the standard normal integral given in Equation 4. RTSYS uses a rational function approximation for Φ that gives greater than single precision accuracy (Kennedy & Gentle, 1980, pp. 90–92).

RTSYS replaces maximization of log likelihood with an equivalent operation, minimization of minus log likelihood. Since $l(RT|\mu, \sigma, \tau)$ is usually negative, $-l(RT|\mu, \sigma, \tau)$ is usually positive and acts like a lack of fit measure, such as squared deviation in least squares fitting: a minimum value of minus log likelihood indicates the best-fitting model.

Search for the Maximum Likelihood Estimate

To illustrate the process of maximum likelihood estimation through search, we will investigate fitting of an example data set of 10 RTs. An ex-Gaussian distribution with $(\mu, \sigma, \tau) = (500, 50, 100)$ was sampled, resulting in the data set (474.688, 506.445, 524.081, 530.672, 530.869, 566.984, 582.311, 582.940, 603.574, 792.358). Since the ex-Gaussian is a convolution, the easiest way to generate samples from it is to sum pairs of samples from normal and exponential distributions. Table 1 gives estimates of the mean and second and third central moments (transformed to the same scale) and the ex-Gaussian parameters obtained by both maximum likelihood estimation and method-of-moments estimation. Equations 1–3 were used to convert method-of-moments estimates of $\mu_1, \mu_2,$ and μ_3 to ex-Gaussian parameter estimates and maximum likelihood estimates of $\mu, \sigma,$ and τ to moments estimates.

The sample has positive skew as indicated by a median (549) less than the mean and has positive estimates of the third central moment. Both methods produce similar estimates. Note that this occurred because the sample was selected to have a positive third central moment estimate. The sample size used is too small for precise estimation and would often result in negatively skewed samples. For example, the method-of-moments third central moment estimate is dominated by the cubed residual for the largest data point that is an order of magnitude greater than all other cubed residuals. When this point is removed, the method-of-moments estimate of $\mu_3^{1/3} = -54.1$, indicating negative skew, despite the fact that the median (530.9) is still less than the mean (544.7). Ratcliff (1979) criticized the method of moments for its sensitivity to extreme RTs, which may be outliers not

Table 1
Method of Moments and Maximum Likelihood Estimates
(Assuming an Ex-Gaussian Distribution) of the First Three
Central Moments and Ex-Gaussian Parameters
for an Example Data Set

	$\hat{\mu}_1$	$\sqrt{\hat{\mu}_2}$	$\sqrt[3]{\hat{\mu}_3}$	$\hat{\mu}$	$\hat{\sigma}$	$\hat{\tau}$
Method of moments	569.5	87.7	102.9	487.8	32.0	81.7
Maximum likelihood	569.5	77.5	92.9	495.8	24.0	73.7

generated by the psychological process under investigation. In contrast, the ex-Gaussian-based maximum likelihood estimate remains positive ($\mu_3^{1/3} = 11.9$) when the largest value is dropped from the sample.

The process of search used to obtain the maximum likelihood estimates can be illustrated geometrically. Assume that we know the true value of two of the three ex-Gaussian parameters. We could then construct a plot of minus log likelihood as a function of the unknown parameter. Figures 3a, 3b, and 3c illustrate such plots for the example data when the unknown parameters are μ , σ , and τ , respectively. The value of the unknown parameter with the lowest value on the plot is the maximum likelihood estimate of the parameter.

Similarly, if we knew the true value of only one parameter, we could plot each value of the unknown pair of parameters on a plane and the corresponding minus log likelihood as a point on a surface above the plane. Figures 3d, 3e, and 3f plot the surfaces for μ , σ , and τ known, respectively. The maximum likelihood estimate corresponds to the coordinates of the point on the plane under the lowest point on the surface. The likelihood surface for estimation of all parameters simultaneously cannot be illustrated because it requires four dimensions (three for the parameters and one for the likelihood).

Determination of the entire likelihood surface is wasteful (computation of the three dimensional plots took several hours each, see Appendix) when we require knowledge of only one point, the minimum. Fortunately, in many cases, we can find the minimum through search with only local knowledge of the surface. To illustrate, imagine that you wish to find the lowest point in a land-

scape, but it is dark so you have only knowledge of the local height provided by moving your foot around. You could proceed by determining the lowest point in your neighborhood, moving to it, then repeating the process until you are at the minimum and are surrounded by higher points.

The problem of search or optimization has been intensively studied, and algorithms exist that make search tractable even in high dimensional spaces. RTSYS uses two such algorithms: simplex and Marquardt-Levenberg. Press, Flannery, Teukolsky, and Vetterling (1988) give detailed explanations and the exact algorithms. The remainder of this section gives a brief description sufficient to facilitate use of the algorithms within RTSYS.

Simplex is a simple and robust algorithm that acts very much like the above analogy, except that it moves around by contracting, expanding, and moving a four-dimensional hypercube (the simplex). The Marquardt-Levenberg algorithm is more sophisticated, using the first and second derivatives of log likelihood with respect to the parameters as well as its magnitude. The first derivatives (slopes) are used to determine the direction of steepest descent, and the second derivatives are used to estimate how far to step assuming that the surface is quadratic (e.g., a good approximation for Figure 3a). Search is terminated when fractional tolerance, the ratio of the previous improvement in likelihood to the improvement of likelihood on the present move, falls below a criterion value. Note that termination is not guaranteed. When the surface around the minimum is relatively smooth, termination occurs rapidly. However, when the surface is irregular, termination may be slow or may not occur at all. Generally, larger RT

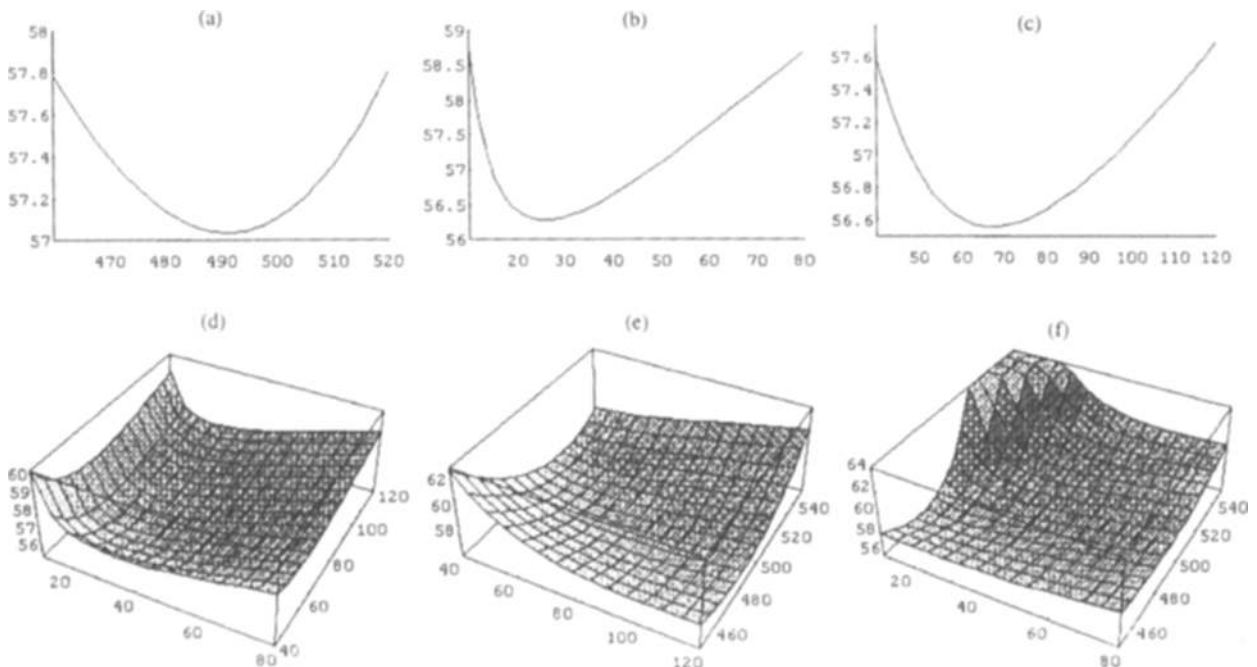


Figure 3. Cross-sections of minus natural log likelihood surfaces, where likelihood is on the vertical axis, with ex-Gaussian parameters (μ, σ, τ) held constant at (500, 50, 100) except for (a) $\mu = 460-520$, (b) $\sigma = 10-80$, (c) $\tau = 40-120$, (d) $\sigma = 10-80$ and $\tau = 40-120$, (e) $\mu = 450-540$ and $\tau = 40-120$, (f) $\mu = 450-540$ and $\sigma = 10-80$.

data sets result in smoother likelihood surfaces and termination in fewer steps.

Under ideal conditions, the Marquardt-Levenberg algorithm will find the minimum much more quickly than will the simplex algorithm. However, certain surfaces, such as a gradually descending valley, can cause very slow convergence. Heading in the direction of steepest descent, combined with some overshoot of the bottom of the valley, may lead to repeated crossing at almost right angles to the gradual decrease toward the minimum. The surface may also be markedly nonquadratic, leading to poor estimates of how far to move on each move on step.

Experience with both methods indicates that the simplex algorithm is the most robust, rarely failing to converge. With smaller data sets ($n < 300$), the Marquardt-Levenberg algorithm tends to move quickly to the region of the minimum but then either does not converge or converges slowly. For larger data sets, however, the Marquardt-Levenberg method is much faster than the simplex method. For intermediate cases, RTSYS allows a combination of the two methods, using the Marquardt-Levenberg method for a default number of steps to find the approximate region of the minimum, then switching to the simplex method to find its exact location.

Problems With Search

In general, search is a sensitive process that requires careful attention. It often results in nonconvergence or error conditions due to floating point overflows or underflows. With a slow computer, or even with a fast computer and a large number of data sets, determining the best-fitting ex-Gaussian parameters may become prohibitively time consuming. The following describes features of RTSYS that automatically take care of many difficulties arising during search. The design goal for RTSYS was to allow unattended computation of best-fitting parameters so that even users of slow machines can process large data sets by leaving RTSYS running overnight.

The geometric analogy illustrates one difficulty with optimization: the problem of local minima. A local minimum is a point that is surrounded by higher points but is still not the overall lowest point, the global minimum (see Figure 4). Whether search descends to a local or global minimum depends upon where it starts. Local minima can be avoided by a judicious choice of the starting point for search.

RTSYS allows the user to select a starting point, but it is also able to estimate its own starting point. When the method of moments produces reasonable estimates of the ex-Gaussian parameters (i.e., $\sigma > 0$, $\tau > 0$; e.g., Table 1), they can be used as a starting point for search. When they are not reasonable, RTSYS uses a heuristic: $\tau = 0.8 \times$ sample standard deviation. The heuristic is based on Luce's (1986) observation previously cited (see also Ratcliff, 1993, p. 510) and my own experience that τ dominates the variability of the ex-Gaussian. The remaining parameters are determined using Equations 1 and 2 (i.e., $\sigma = 0.6\tau$, and $\mu = \text{sample mean} - \tau$).




Figure 4. A two-dimensional likelihood surface with a local minimum and a global minimum.

Some data sets may cause the search algorithm to attempt to evaluate the log likelihood function for $\sigma \leq 0$ or $\tau \leq 0$, causing a floating point error. The ex-Gaussian distribution may be an inappropriate model when the data's distribution is symmetric or negatively skewed ($\tau \leq 0$) or when it is sharply peaked and better modeled by an exponential plus a constant ($\sigma \leq 0$). Such cases may also occur due to sampling error when the ex-Gaussian model is appropriate, especially for small sample sizes. Even when the final result is positive values of σ and τ , attempts may be made during search to evaluate the log likelihood function at parameter values that cause floating point errors. To avoid floating point errors, RTSYS sets the minus log likelihood function to a large value for σ or $\tau \leq (\text{sample mean})/1,000$. This effectively introduces a barrier in the likelihood surface ensuring that search does not move into inappropriate regions. Failed evaluations of the likelihood function are reported by RTSYS and implausible parameter estimates are set to the missing value, alerting the user to the problem.¹

When fitting fails, RTSYS allows the cause to be determined and a remedy, if possible, to be found. The problematic RT distributions can be plotted to examine the shape of the distribution. Search may be attempted again with new start points or other parameters controlling the search process, such as the criterion for terminating search. In particular, simplex search (the most robust method) can be carried out in an interactive mode. During interactive mode, the parameter vector (μ, σ, τ) and minus log likelihood are displayed for each step. At any step, the search may be paused and restarted with new control parameters supplied by the user. Even with these precautions, some cases may not converge (i.e., the search may converge in an unacceptable region of the parameter space or get caught in a limit cycle or even chaotic oscillation). In such cases, censoring the data set or collecting more data may be necessary.

Even when search is not problematic, it is important to ensure that the ex-Gaussian distribution provides an adequate model of the data. RTSYS allows evaluation of the ex-Gaussian model, both graphically, through plotting the ex-Gaussian curve on a histogram of the data (see Figure 5) and, inferentially, through a chi-square test.




Figure 5. The histogram of RT frequency given in Figure 1, with the best-fitting ex-Gaussian function (solid line) superimposed.

Chi-square is calculated by comparing the observed and expected number of RTs in each of a series of categories that span the range of the RT distribution.

A difficulty with chi-square testing is choosing the width of the categories, especially when the scale of distributions may vary widely. D'Angostino and Stephens (1986, p. 69) recommend the use of chi-square cells with equal probabilities under the fitted distribution, citing a reduction in bias and better small sample properties. Unfortunately, determination of the width of equal probability bins requires search using the integral of the ex-Gaussian pdf, an operation more computationally expensive than parameter fitting itself. Instead, RTSYS uses categories with equal numbers of data points, approximating equally likely ex-Gaussian categories for reasonable fits. The user may select the number of categories or allow RTSYS to automatically select $2n^{2/5}$ categories (a heuristic suggested by D'Angostino & Stephens, 1986, p. 70). When a category produces less than five expected or observed values, RTSYS collapses the category with the subsequent category to ensure that the assumptions of chi-square testing are not violated.

A Nonparametric Measure of Distribution Asymmetry

A nonparametric measure of distribution asymmetry is provided by the standardized distance between the mean and median [i.e., $(\text{mean} - \text{median})/\text{standard deviation}$]. The measure is briefly discussed by Kendall, Stuart, and Ord (1987, p. 106, and Exercise 3.22, p. 116) and was used by Ratcliff (1993). It is bounded between -1 and 1 , with positive values indicating that the mean is greater than the median, as is commonly observed for RT distributions. Ulrich and Miller (1994) point out that this measure should be less than 0.31 (the theoretical value for an exponential distribution) for RT distributions.

A related measure, which replaces the median with the mode, has been studied in relation to the unimodal Pearson system of distributions (see Kendall et al., 1987). However, RTSYS calculates the median version because the mode is difficult to estimate efficiently. Although values of the median form of the statistic, called *nonparametric skew* in RTSYS, cannot be used to produce estimates of the third central moment and hence estimates of parametric skew, they provide an alternative way of assessing skew in situations where the sample size is small. Nonparametric skew's advantage over the method-of-moments estimate is that it is not as sensitive to outliers.

Summary

Estimation of skew using the method of moments has unacceptably high variance for sample sizes commonly used in experimental psychology. RTSYS provides a more efficient way to determine skew through maximum likelihood estimation of the ex-Gaussian distribution. The researcher can choose either robust simplex search with smaller samples or more efficient but less robust Marquardt-Levenberg search with large samples (such as may be produced by simulations; e.g., Mewhort et al., 1992) and assess the suitability of the ex-Gaussian model with chi-square testing. For samples that are too small to allow fitting of the ex-Gaussian (a minimum of approximately 100 observations is recommended), RTSYS can calculate a nonparametric measure of distribution asymmetry. Log and inverse transformations are also provided for use with small samples and inferential tests that assume normality. An alternate solution to the problem of sparse data is given in the next section.

Vincentizing: What to Do When Data Are Sparse

The most frequently encountered difficulty in dealing with skew arises from small data sets. Even maximum likelihood estimates of skew are less efficient than the sam-

ple mean. Hence, data sets sufficient to make reliable inference about the mean will not always be adequate for reliable inference about skew. As samples become smaller sampling error leads to increasingly frequent missing values for skew estimates, further complicating inference. Typically, experiments that attempt to measure skew have at least 100 observations; however, in many paradigms, 100 observations is an unrealistic goal. Difficulties range from limited stimuli, limited time, and limited attention on the part of subjects to confounding by practice effects. The usual strategy adopted by the behavioral scientist when faced with low numbers of observations, and consequent variability in estimates, is pooling or mixing of observations from different conditions or subjects. When measuring the shape of distributions, however, mixing may lead to artifacts. The artifacts can be illustrated with a simple example using the familiar Gaussian or normal distribution.

Suppose two experimental units (e.g., subjects or conditions) produce normally distributed data with different means and variances (Figure 6a). Mixing the data over units results in an almost bimodal rather than normal distribution (Figure 6b). Ideally, we want a combination technique that preserves the component distribution's shape while averaging their parameters (Figure 6c). However, when we do not know the underlying distribution (e.g., RT data), or when we assume the distribution but cannot estimate its parameters with few observations (e.g., fitting the ex-Gaussian fails), we cannot average parameters.

Ratcliff (1979) proposed that Vincent averaging, a technique originally used to average learning curves (Hilgard, 1938; Vincent, 1912), provides the stabilizing effect of averaging without distortion of distribution shape. He demonstrated analytically that Vincent averaging produces no shape distortion for several theoretical distributions (exponential, logistic, and Weibull), and he demonstrated, using Monte Carlo techniques, that distortion is minimal for the ex-Gaussian. He also analyzed a large RT database with sufficient observations for stable fitting of the ex-Gaussian distribution to individual subject's data. Ex-Gaussian parameters found by fitting to Vincent-averaged distributions were close to the average of ex-Gaussian parameters estimated from individual subject's data. Hence, Vincentizing was able to produce

average distributions with the same shape as component distributions and parameters that were the average of component distributions' parameters.

Vincent averaging is achieved in two stages. First, quantiles are calculated for each component distribution. If, for instance, five quantiles are required, calculate the times below which 10%, 30%, 50%, 70%, and 90% of observations fall for each component (e.g., subject). Second, corresponding quantiles are averaged across components. The intervals between adjacent averaged quantiles contain equal proportions ($1/q$, where q is the number of quantiles) of the average distribution's RTs. The intervals below and above the first and final quantiles each contain $1/(2q)$ of the RTs.

Ratcliff (1979) suggested a simple algorithm for calculating component quantiles. For n ordered reaction times and q quantiles, write an array containing q copies of each of the n RTs in turn. If, for instance, the RTs (1,2,3) are observed and two quantiles are required, then the array contains (1,1,2,2,3,3). Quantiles are obtained by averaging successive groups of n observations from the array [e.g., $(1+1+2)/3 = 1.33$, and $(2+3+3)/3 = 2.66$]. Dawson (1988) provides an implementation of the algorithm. RTSYS uses a modified version of Ratcliff's algorithm that addresses a problem arising with large data sets. RTSYS can process a maximum of 2,500 RTs and 100 quantiles. With each RT being a 6-byte real, a direct implementation of the algorithm requires up to 1.5 MB of memory! RTSYS avoids high memory loads with an algorithm that refills the same n element array after each quantile is calculated. The slight bookkeeping overhead is far outweighed by reduced memory requirements.

The ex-Gaussian distribution can be fit directly to the average quantiles because they form a representative sample from the average RT distribution (i.e., by definition, each average quantile RT is equally likely to be sampled). An added advantage of Vincentizing is increased speed in fitting the ex-Gaussian, at least with data sets where $q \ll n$. The speedup occurs because calculation of likelihood requires the pdf for individual observations to be evaluated q times rather than n times.

For good ex-Gaussian fits, Vincentizing requires as few as 20 observations per condition, with the restriction that $q < n$ (and preferably, $q < n/3$). When small numbers of ob-




Figure 6. (a) Two normal distributions with means of 0 and 3 and variances of 1 and 4, respectively. (b) The distribution resulting from pooling or mixing the two normal distributions. (c) The distribution resulting from averaging the component parameters (mean = 1.5, variance = 2.5).

servations per condition are used, Vincentizing should be performed over more components to ensure a stable estimate. Generally, q is chosen large enough so that the shape of the distribution is preserved but small enough so that high-frequency noise is not included (usually $5 \leq q \leq 20$).

The choice of number of quantiles is aided by plotting a Vincent histogram (see Figure 7). A Vincent histogram is a graphical representation of the average distribution. It is plotted as $q - 1$ rectangles, which span the intervals between adjacent average quantiles and with heights chosen so that their areas are $(1/q)$ (Ratcliff, 1979). The Vincent histogram is a discrete approximation to a pdf. Like a pdf, the area under the Vincent histogram in a range is the probability of sampling an RT in that range. The shape of the Vincent histogram can be used to judge whether enough quantiles have been chosen to provide a smooth estimate of the pdf without blurring its shape. The best-fitting ex-Gaussian function can be superimposed on the Vincent histogram to determine how well it models the data. The suitability of the ex-Gaussian model for Vincentized data can also be assessed inferentially with a chi-square test.

Standard Errors for Vincent Averages

A drawback of Vincentizing over subjects is that it removes subject variation. Hence, inferential testing on the parameters of the fitted ex-Gaussian cannot be performed. Ratcliff and Murdock (1976, p. 200) provide a table of standard deviation estimates for a range of ex-Gaussian parameter values. The table is derived using large sample properties of maximum likelihood estimators and may not be appropriate for small samples and parameter estimates derived by Vincentizing.

A more satisfactory solution for small samples may be provided by resampling (e.g., Efron & Tibshirani, 1993). Resampling involves creating a number (e.g., 100) of new Vincent averages from the original data sets by sampling with replacement (i.e., select a subject at random,

then repeat the process until n subjects are chosen, then calculate the Vincent average of the sampled data sets). The ex-Gaussian distribution is then fit to each of the resampled Vincent averages. The standard deviations of the parameter estimates can be used to perform inference.

In the current version of RTSYS, resampled data sets must be created externally then processed in the usual way by RTSYS, most conveniently with the resampling replications as a factor. At present, I am investigating the theoretical properties of resampling combined with Vincentizing. Future versions of RTSYS will incorporate automatic resampling if the results are favorable.

Summary

Vincent averaging allows RTSYS to determine skew even when as few as 20 observations are available for each experimental unit in each cell of the design. Vincent averaging combines data across subjects or conditions without distorting the shape of the underlying distribution. The ex-Gaussian distribution can then be fit to the Vincentized data to obtain a maximum likelihood estimate of skew. A chi-square test allows assessment of the adequacy of the ex-Gaussian as a model of the data. The distribution of the Vincent-averaged data can be viewed by plotting a Vincent histogram. The Vincent histogram also provides a visual summary of the average distribution of the data, which is useful even when large numbers of observations are available for each cell.

Censoring

Typically, RT data are used to investigate a particular psychological process. The process of interest may not, however, be the only process to contribute to RT variation. Some confounds can be controlled by careful experimental design, but two confounds generic to RT data—anticipation and distraction—may require post hoc measures. Anticipation produces fast RTs, and distraction




Figure 7. A Vincent histogram with 20 quantiles. Each rectangle has an area of $1/20$, and there are 19 rectangles.

produces slow RTs. An indirect solution is to use a robust measure of central tendency, such as the median, which is insensitive to extreme observations. This does not, however, solve the problem when variance and skew are of interest. A direct solution is to remove or censor uncharacteristically fast and slow RTs. Unfortunately, a poor choice of censoring criteria can distort the nature of the process of interest by removing real observations. The positive skew of RT distribution makes censoring criteria for slow RTs particularly problematic (see Ratcliff, 1993, and Ulrich & Miller, 1994, for extensive reviews of censoring methods for RT data).

RTSYS supports three types of censoring criteria: absolute, standard deviation, and percentile. Censoring criteria can be specified separately for fast and slow RTs, allowing the asymmetry of RT distribution to be taken into account. Absolute censoring allows specification of times above or below which observations are rejected. Like all censoring, it should be used with caution because it may introduce bias. Ulrich and Miller (1994) study absolute censoring extensively and provide a useful set of recommendations for how to safely use it. Absolute censoring is most appropriate for rejecting fast RTs that occur through anticipation. Experiments on simple RT to intense signals indicate that the residual (nondecision) component of RT is at least 100 msec (Luce, 1986, Table 2.1, p. 62). Hence, routine use of an absolute fast RT censoring criterion of 100 msec seems justified.

Standard deviation censoring rejects RTs that are a criterion number of standard deviations above or below the mean. Percentile censoring rejects RTs above or below a criterion percentile point. Percentile and standard deviation censoring techniques are useful because they adapt to the scale of the data's distribution. For instance, when comparing cells with different means, absolute censoring will affect means and proportions of observations censored differently, biasing comparisons between conditions. Standard deviation censoring will adapt to different condition means, and percentile censoring will ensure that equal numbers of observations are removed in each condition.

Percentile censoring should be treated with care when comparing conditions in which outliers occur with different probabilities. For example, an irritating or tedious condition is more likely to produce slow outliers due to distraction than is an enjoyable condition. In such a case, where it is inappropriate to remove the same percentage of slow data from both conditions, standard deviation censoring may be preferred. RTSYS also calculates mean error RT so the predominant type of error (faster or slower than correct mean RT) can be used to guide criterion setting (see Link, 1982, and Pachella, 1974, for a discussion of speed-accuracy tradeoff).

In a study of the effect of outliers on the power of analysis of variance, Ratcliff (1993) recommends either an inverse transformation when variability among subject

means is low or standard deviation censoring when subject variation is high. He also suggests that the researcher try "a range of cutoffs and make sure that an effect is significant over some range of nonextreme cutoffs" (Ratcliff, 1993, p. 519). Censoring and rescaling facilities in RTSYS make it easy to implement his recommendation.

Compensating for Censoring With Maximum Likelihood Estimation

Even when censoring is informed by the experimenter's expertise, it is still likely to reject some real data. Fortunately, with samples large enough to fit a distribution, maximum likelihood estimation can be used to compensate for the unwanted effects of rejection of real data. As when determining skew, the maximum likelihood technique requires assumption of a theoretical distribution. The effect of censoring r_1 fast and r_2 slow observations from a sample of n RTs is incorporated into the likelihood function, L , of a theoretical distribution with pdf $pd(RT)$ is

$$L = \left(\int_{-\infty}^a pd(RT)dt \right)^{r_1} \times \prod_{i=1}^{n-r_1-r_2} pd(RT_i) \times \left(\int_b^{\infty} pd(RT)dt \right)^{r_2} \quad (5)$$

(note that Π indicates a repeated product). The integrals assess the probability of an RT less than a and greater than b , where a and b are the lower and upper censoring criteria. Hence, the first and last terms assess the probability of observing r_1 RTs less than a and r_2 RTs greater than b . The middle term is the likelihood of the $(n - r_1 - r_2)$ uncensored observations. Although other adaptations of the likelihood function are possible (cf. Kendall & Stuart, 1967, Equation 32.36, p. 523), Ulrich and Miller (1994) found fitting with Equation 5 to be the most robust approach.

Ulrich and Miller (1994) carried out a set of simulations to determine how fitting with Equation 5 affects bias in estimating means and standard deviations. They found good bias reduction properties with little decrease in efficiency, especially for mean estimates. They recommend an Erlangian distribution for best recovery of the mean and the ex-Gaussian for best recovery of standard deviation. However, there is little to choose between the Erlangian and ex-Gaussian for mean estimation (average relative biases ranging from -0.1 to 2.7 and from -1.5 to 1.4 , and maximum inefficiencies of 1.34 and 1.81 , respectively; see Ulrich & Miller, 1994, Tables 14 and 15, p. 68-69). RTSYS assumes an ex-Gaussian pdf in Equation 5 and automatically fills in the values of r_1 , r_2 , a , and b implied by the censoring criteria selected. Hence, RTSYS provides a convenient mechanism for implementing the approach recommended by Ulrich and Miller.

Summary

RTSYS provides three criteria for censoring data: absolute, standard deviation, and percentile. Criteria can be applied independently and in any combination for both slow and fast RTs. The criteria are applied uniformly for each cell of the experimental design. The ex-Gaussian likelihood function can be automatically adjusted depending on the number of observations censored in each cell. This procedure entails increased computation,² but it reduces bias introduced by censoring real RTs along with outlier RTs.

THE USER INTERFACE

The RTSYS user interface consists of a Main menu and six submenus. Each menu has options that can be selected either by pressing the capitalized letter of the option or by moving a highlighted bar with the cursor keys and pressing enter when the desired option is highlighted. Context-sensitive help can be obtained by pressing the F1 key when the highlighted bar is on the option in question.

Main Menu

The Main menu consists of options and status indicators. Status indicators appear at the bottom of the screen and, if defined, include the names of active files. The Main menu options are submenus (except for Help and Quit) that have a similar format to the Main menu, but without the status indicators. Submenus consist of options on the left of the screen and default settings on the right of the screen. The defaults control the way in which the options are performed. Selecting the Defaults option of the submenu causes the Defaults menu to appear. Default values can then be changed, by either entering explicit values or toggling between possible states of the default.

Defaults can also be changed through the Defaults submenu of the Main menu. Defaults set in this way are changed permanently through storage in a defaults file. RTSYS reads a defaults file on start up and is distributed with a standard defaults file, *rtsys.def*, in its home directory. Defaults files can also be stored by the user in data directories using the Save option of the Defaults submenu. This should be done to customize RTSYS when performing a series of similar analyses. When RTSYS is executed using the *rt.bat* file (see next section) from a directory containing a defaults file, it will automatically load that defaults file, allowing analysis to proceed with appropriate settings.

Most options require information about the design of the experiment that you are analyzing. Design information is specified in the Files menu and stored in a design file. Options may also require that RTSYS has information about file names. If a design file and file names are not specified, RTSYS will ask appropriate questions after you select an option. You may not always know the answers

to the questions, especially if file names are requested. If this happens, you can simply hit enter and you will be returned to the previous level.

RTSYS submenus are used in a sequence to process data files and output results. In order to speed later calculation of statistics, experimental data files must be combined with design information to create ordered files. Ordered files are ASCII and are created and manipulated in the Data Management menu. Vincentizing can also be performed in the Data Management menu. Once ordered files are created, frequency or Vincent histograms can be plotted in the Graphics menu to check the data's distribution. Statistics are calculated from ordered files and stored as (non-ASCII) statistics files in the Analysis menu. Once the ex-Gaussian parameters are calculated in the Analysis menu, the Graphics menu can be used to view the best-fitting distribution superimposed on histograms. The contents of statistics files can be viewed, printed, and saved to ASCII parameter files in the View Statistics menu. While both Data Management and Analysis menus create one file corresponding to each data file, the View Statistics menu creates, for each statistic, a single parameter file that combines results from between-subject and within-subject design cells. The following sections provide an overview of the options and defaults of each submenu.

Specifying File Names and Interacting With DOS: The Files Menu

To make finding out about files easier, the Files submenu is available from the Main menu and all submenus. To facilitate single-task usage under DOS, the Files menu allows you to interact with the operating system. Options provided allow you to list files (the DOS "dir" command), look at ASCII files (the DOS "type 'file name' | more" command), edit ASCII files (using the edlin editor), delete, rename, and copy files (the DOS del, ren, and copy commands), and create and remove directories (the DOS md and rd commands). Under Windows, most of these functions are more easily accomplished with the File Manager. The remaining Files menu options allow you to specify file names and design information.

RTSYS can be used in both *single-file* and *multifile* modes. In single-file mode, operations are performed on a single file. If a file required by an option is not defined, the user will be asked to specify it on a command line. The names of each type of file (Data, Ordered, or Statistics files) can also be specified using options in the Files menu. In multifile mode, all operations are done on a set of files whose names are specified in a holding file. The holding file is an ASCII file containing one file name per line. A holding file is specified using the Multiple Files option in the Files menu. Holding files allow large sets of files to be batch processed and allow specification of subjects and between-subject factors. It is best to specify only the root names of files in the hold-

ing file. In multifile mode, RTSYS attempts to append file extensions for data, ordered, and statistics files as required by the option. This allows a holding file listing only the root names of the data files to be used to execute several different options without having to respecify the names on each occasion. Default extensions used are defined in the Files Defaults menu.

A within-subject design must be specified by the user and stored in a design file. Each line of the design file contains the levels of each factor for a cell. Levels are specified by integer designators. You must specify factors in the same order as they occur in your data files and use the same designators as used in your data files. Design files can be created from within RTSYS using the Get Design File option in the Files menu. The Get Design File option generates factor level indicators 1, 2, 3, and so on, by default, but it allows recoding to more meaningful (unique integer but not necessarily ordered) values if required.

This method of creating design files works only for fully factorial designs. If you have a nonfactorial design, you can create a design file either directly using an ASCII editor (access to DOS's `edlin` is given in the Files menu) or by creating the full factorial design of which your design is a subcase and then editing out the cells (lines) that are not present.

Once a design file is created, it can be loaded from disk using the same option. Design files control the order in which cells are processed. This is particularly useful when generating parameter files for input to other applications. Design files can also be used to ignore cells by omitting the lines that designate them. Operations in all submenus are performed only on cells specified in the design file.

All Files options assume the Default Path, which you can set in the Files Defaults menu, unless you begin your specification with a "\". For example, if you set the Default Path as "\dat\expl", RTSYS will search for a file specified by *new.dat*, as "\dat\expl\new.dat".

Data File Format: The Data Management Menu

Two types of data files are accepted by RTSYS: blocked and randomized. Blocked files use only one condition or cell of the within-subject design in a block of trials. Randomized data files come from a design in which different cells are mixed in a block.

Blocked Data Files

In a blocked data file, RTs from each block are preceded by a cell designator line. The cell designator line begins with a block header indicator (an integer value that cannot equal an RT value), followed by integer values specifying the levels of each factor for a cell. Following lines contain an integer correct/incorrect/discard trial indicator (in the default state, 1 = correct, 2 = incorrect, 3 = discard trial) and an RT value (which can be real or integer). For example,

```
-3 5 3 2 {Cell designator line, Block header (-3),
          cell (5 3 2)}
1 678   {Correct RT}
1 735   {Correct RT}
2 1111  {Error RT}
...
```

Note that comments in braces { } are not in the actual file.

If the data are not of the accuracy type (i.e., there are no correct and incorrect answers), the trial indicators before the RT can be omitted. The type of data file assumed (accuracy or not) is controlled by the Accuracy Data default in the Data Management Defaults menu.

Randomized Data Files

Each line in a randomized data file contains integer values designating the levels of each factor for a trial, a trial indicator, and an RT value (real or integer). For example,

```
5 3 2 1 679 {Cell 5,3,2 Correct RT}
3 2 2 1 699 {Cell 3,2,2 Correct RT}
4 1 1 2 799 {Cell 4,1,1 Error RT}
...
```

The type of data file to be used and the values of the different indicators can be set in the Defaults section of Data Management menu. Again, assumption of either accuracy (final value before the RT is a trial indicator) or nonaccuracy data formats is controlled by the Accuracy Data default in the Data Management Defaults menu.

The user may specify that any number of values be ignored before a factor specification (after the Block Header Indicator in a blocked file or at the start of each line in a randomized file) and also before the trial indicator in a blocked file. This may be useful if specific information about a trial, such as a trial number, is included in a data file. The user can also specify collapsing by choosing which factors to ignore. Again, the Defaults section of the Data Management menu can be used to set these values. Finally, header lines in the data file may be skipped using the Data File Header Line Skip option in the same defaults menu. A maximum of 2,500 (correct + error) RTs may be analyzed in each design cell.

RTSYS does only cursory analysis of error data (mean and percent). This is appropriate in most RT paradigms because errors are kept low to avoid speed-accuracy tradeoff. When error rates are low, even estimates of error variance are imprecise, let alone ex-Gaussian parameters. If, however, you want more extensive analysis of error data, the data file can be processed with indicators for correct and error data flipped (usually set in the Files Defaults menu, because this is easier than changing the actual data file) so that RTSYS treats the error data as correct data and treats correct data as error data. Alternatively, you can use correct/incorrect as a factor in the design and disable the Accuracy Data default in the Data Management menu.

Ordered Files

In an ordered file, RTs from each cell are collected together and listed from smallest to largest. Given that most RT data sets consist of less than 1,000 observations, I follow Press et al's. (1988, p. 227) recommendation and use Shell's method of sorting. RTs are also divided according to whether they are from correctly answered or error trials. When creating ordered files, RTSYS allows RTs to be linearly rescaled to change time units. RTs can also be transformed according to the resolution of the experimental apparatus that collected them. Most RT measurement systems assume a measurement unit, usually 1 msec or the frequency of the monitor refresh (16.7 or 20 msec). RTs are actually measured as the number of measurement units accumulated before a response is detected. Hence, the best estimate of the actual RT is one half a measurement unit greater than that actually recorded. Although the difference is usually small, it can be important for order statistics, especially with low-resolution measurements (see Heathcote, 1988, for a discussion of time measurement in computer-controlled experiments).

In an ordered file, RTs for each cell are preceded by a cell designator line (similar to a blocked data file but with no block header indicator). Each line thereafter contains a single RT value. The list of correct RTs is followed by an indicator (-1 as default value), and the error RTs are listed. The end of the error RTs is marked by another indicator (-2 default), and the next cell then begins. For example,

```

1 1      {Cell 1,1}
567
784
733
-1      {End of Correct RTs}
899
-2      {End of error RTs}
2 1      {Cell 2,1}
967
...

```

Ordered files can be produced directly from single data files or by Vincent averaging across a number of ordered files. A holding file specifying the names of the ordered files to be Vincentized over must be supplied (using the Files menu), because file names cannot be entered from the command line. Ordered files can also be merged using the Join Ordered Files option of the Data Management menu. For example, you may wish to join together the data from one subject collected over several days. Again, a holding file specifying the names of the files to be combined is required.

Calculating Statistics: The Analysis Menu

The analysis menu produces statistics files from ordered files. Statistics files contain one record for each cell of the design. Each record contains the statistics described above. Input from the ordered files can be non-

linearly rescaled (using inverse and log transformations) and/or censored. Censoring and rescaling are controlled from the Analysis Defaults menu. They have no permanent effect on the ordered files, only the results stored in the statistics files. Similarly, factors may be collapsed in the Analysis menu and results stored in statistics files with no effect on the original ordered files. For flexibility, it is advisable to specify ordered files with the finest factor structure applicable to your data and collapse nuisance factors during analysis.

You will often wish to try different combinations of censoring, rescaling, and collapsing. In multifile mode, you can use the same holding file for each combination but maintain separate file names for each combination by changing the default Statistics File Extension in the Data Management Defaults menu. For example, you may wish to compare statistics for uncensored data (e.g., stored in *.sts files) with statistics for data with the upper 5% of observations removed (e.g., stored in *.su5).

Analysis can be done with or without fitting the ex-Gaussian distribution function. Not fitting the ex-Gaussian is much quicker and may be appropriate for a first look at the data. If the ex-Gaussian is not fitted, its parameters will be set to the missing value in the statistics files. Although every attempt has been made to make fitting automatic, some data sets may require closer scrutiny. Such scrutiny is aided by some knowledge of the fitting process (see the previous section, Search for the Maximum Likelihood Estimate).

Search times for fitting the ex-Gaussian may be slow, especially on XT or AT machines. In such cases, multifile mode can be used to perform overnight runs. You should experiment with the tolerance parameter in the Analysis Defaults menu to determine what values give you an accurate fit for the least computation. The Tolerance Default determines how small the relative decrease in minus log likelihood must be between steps of the fitting algorithm before it stops. If you can increase the tolerance without changing the estimated parameters, do so; this will speed fitting. Do not set tolerance less than 10E-8, because fitting will converge slowly due to spurious variation introduced by rounding errors.

The search algorithm will abort after a number of iterations specified in the Analysis Defaults menu, and it will continue fitting with the next cell. This default avoids cases where fitting is not converging. Set it to a large value to ensure that a reasonable attempt at fitting is made, but not so large that multifile runs are unable to complete in a reasonable time.

Viewing Results: The Graphics and View Statistics Menus

RTSYS has two menus dedicated to viewing and output of the results of analyses. The Graphics menu is mainly designed to help fitting the ex-Gaussian, although limited options for formatting and output of graphs are available. The View Statistics menu allows viewing of statistics and output to either a printer or a file.

Graphics

The Graphics menu allows you to look at the distribution of your data. Two types of histograms can be plotted: frequency histograms and, for Vincentized data, Vincent probability histograms. Plotting histograms requires only that you have created ordered or Vincentized files. You can specify the scaling of the histograms, using the Graphics Defaults menu. Otherwise, RTSYS will automatically choose the scale to fit the available graphics. If user-specified scaling truncates the data, a warning message will appear and the histogram will not be displayed. User-specified scaling is useful so that direct comparisons can be made between figures.

Once the ex-Gaussian distribution has been fit in the Analysis menu, and the resulting statistics files are available, a plot of the best-fitting ex-Gaussian distribution can be superimposed on either frequency or probability density histograms. Note that you must make sure that ordered and statistics files correspond to obtain correct results!

You can also perform censoring in the Graphics menu using the Graphical Censor option. Unlike censoring in the Analysis menu, graphical censoring creates a new copy of the ordered file (with extension “.cns” in multifile mode). For each cell, histograms and lists of fast and slow RTs are displayed, and you will be asked to supply a fast and slow RT censoring criterion. The histogram is then redisplayed with the censored data removed, and the cycle may be repeated as required. Viewing the histogram gives a sense of the scale of the RT distribution and can be particularly useful for removing outlying RTs due to recording errors.

There are no direct facilities for exporting figures (e.g., the Windows clipboard is not supported), although some memory resident applications are available that directly dump a graphics screen to the printer. Fortunately, many other applications allow you to plot a histogram using the contents of an ordered file. However, it may not be easy to plot the ex-Gaussian density based on the fitted parameter values. To circumvent this problem, RTSYS can save ordered files containing the value of the ex-Gaussian density corresponding to each RT observation. The saved file is identical to an ordered file except that each RT value is followed by a tab character and the corresponding ex-Gaussian value. The two columns of RT and ex-Gaussian values can be used in a graphics package to plot the ex-Gaussian function. It is also difficult to plot a Vincent histogram. RTSYS allows you to save files of coordinates for the corners of the Vincent histogram. The file of tab-delimited coordinates can then be used to plot the Vincent probability histogram with a graphics package.

Note that appropriate Turbo Pascal font files (*.chr) and graphics driver files (*.bgi) must be resident in the RTSYS home directory if you invoke any of the options using graphics. These files should be automatically copied during installation.

Viewing Statistics

The contents of statistics files can be viewed on the screen, printed, or saved in an ASCII file. The output des-

tinuation is determined by a default of the View Statistics menu. In multifile mode, averages across statistics from the same cell in different files are calculated. The output format (including number of columns, decimal places, missing value string, and delimiter for output to file) can be controlled from the View Statistics Defaults menu.

Files appropriate for inferential statistical analysis can be generated in multifile mode. The ASCII file produced, called a *parameter file*, is a listing of values of a statistic for each cell for each subject. The order of subjects (files) is determined by the holding file, usually with one subject's statistics per line. Within-subject statistics are listed within a line, with cell order determined by the design file and values separated by either a space or tab delimiter.

Multifile operations in the View Statistics menu can be slow with large factorial designs due to disk access time. Greater speed can be obtained using the Multifile Speed default. Usually, RTSYS resets a statistics file and searches from the beginning for each design cell. This allows the order of cell output to differ from the order of storage and allows for different storage orders across different statistics files. However, when the Multifile Speed: Fast (No Match) Default is selected, all statistics files are assumed to store cells in the same order as the design file. Hence, statistics files do not have to be reset for each cell, resulting in a faster operation.

Hardware and Installation

RTSYS is distributed on a single 5.25- or 3.5-in. floppy disk. It can be installed to a hard disk by typing *a:\install* from the DOS prompt or using *run a:\install* in the Windows program or File Manager. RTSYS will be installed in the directory *c:\rtsys*, and example data will be stored in the directory *c:\rtsys\testdat*. The *rt.bat* file contained in *c:\rtsys* should be copied into a directory in your system path (e.g., *c:\dos*) so that you may invoke RTSYS from any directory. From Windows, either make a program item that executes *c:\rtsys\rtsys.exe* with the working directory *c:\rtsys* or choose *run rt* from the Windows file manager. The latter method has the advantage that the File Manager working directory will be passed to RTSYS's Default Path. In this way, you can avoid having to type in the path to your data directory. Under DOS, you will be automatically moved to the directory specified in your Default Path when you exit RTSYS. Hence, using the *rt.bat* file under DOS, you can move in and out of RTSYS without having to leave the directory containing your data. You may put RTSYS in a directory other than *c:\rtsys* if you wish. Only the *rt.bat* file and/or the program item directories need be changed.

RTSYS was originally designed for use with slower (e.g., XT or AT) DOS machines. The range of DOS functions available in RTSYS allows single-task usage on a non-Windows DOS machine. Because distribution fitting can be time consuming on slow hardware, RTSYS was designed for robust performance with overnight runs on large data sets. With faster machines (e.g., 486), fitting is relatively rapid. However, disk IO with complex designs may still be a bottleneck because RTSYS holds

information for only one design cell at a time in RAM to avoid exceeding the 64K data segment limit. If extra memory is available, a working directory on a RAM disk greatly speeds file manipulation, especially for large factorial designs. RTSYS 1.0 is configured with a maximum of 299 design cells and 2,500 observations per cell. If you routinely work with either more cells or more observations per cell, recompiled versions are available on request. However, the two settings trade off; therefore, exceeding both settings is not possible.

Flexible parameter file formatting means that RTSYS works easily with most applications. Tab-delimited parameter files can be read into an ASCII editor (e.g., the Notepad) and cut and pasted through the Windows clipboard to most spreadsheets. I routinely use RTSYS in conjunction with spreadsheets in Minitab 9.2 and 10.0 and Word 2.0 and 6.0.

FUTURE DIRECTIONS AND AVAILABILITY

The area of density estimation has undergone significant growth in recent years. Many of the techniques developed are beginning to emerge from the statistical research journals and find practical application. Traditional parametric-model-based approaches, such as maximum likelihood estimation of the ex-Gaussian, are being supplemented by nonparametric techniques that make fewer assumptions about the data. For example, Silverman (1986) reviews kernel-based methods that produce smoothed estimates of density. Kernel methods use local, weighted averaging and are particularly suited for graphical display (the traditional histogram is actually a type of kernel estimate using a rather nonoptimal rectangular kernel applied to disjoint ranges). Tarter and Lock (1993) review another approach, Fourier analysis, and its relationship to kernel techniques. They point out that Fourier methods have an advantage over traditional techniques that use distributions based on elementary functions (such as $\sin(kx)$, x^k , and e^{kx}) because they use a generalized representation, capable of representing any reasonably well-behaved function. Greater generality allows fewer assumptions to be made about the data.

It should not be inferred that "nonparametric" techniques require no assumptions. Both Fourier and kernel methods require specification of the smoothness of the density estimate, a restriction on the complexity of the density akin to the Ockham's Razor heuristic successfully applied in most branches of science. In the Fourier approach, smoothness is determined by the shape of a window on the frequency spectrum. In the kernel approach, it is usually determined by a parameter controlling the width of the kernel. However, these assumptions are preferable to stronger assumptions made by traditional "parametric" approaches, unless, of course, the parametric assumptions are true.

In the future, it is hoped that a selection of these new techniques suitable for RT data will be implemented in RTSYS. In particular, the Adaptive Epanechnikov Ker-

nel seems suited to estimation of RT densities with long tails. As each RT is paired with a density estimate, results of the analysis can be stored as an addition to an ordered file in the same way that ex-Gaussian density estimates can be stored in an ordered file. Appropriately chosen kernel estimates are superior to histograms as estimates of density. Hence, the graphics options can be extended to plot the kernel density estimate, allowing visual checking of the parametric estimate in different regions (e.g., does a parametric tail estimate overestimate or underestimate the tail density?). Nonparametric density estimates may also be useful in recovering estimates of the mode, by examining the derivative of the density estimate, and the median, mean, and especially higher order central moments, through integration of the density. Further investigation is required to determine the usefulness of these techniques.

Accurate nonparametric density estimation is also important for determining the hazard function. Luce (1986) suggested that the tail of the hazard function is useful in discriminating between parametric models with similar densities. The Adaptive Epanechnikov Kernel has been used for hazard function estimation of empirical (Smith, 1995) and simulation (Heathcote, 1995b) data. However, Van Zandt and Ratcliff (1993) point out that hazard function estimation is sensitive to contamination from outliers and mixtures in general, especially in the critical tail region (see also Heathcote, 1995a). Detection and modeling of mixtures is a difficult problem, except in extreme cases resulting in clear multimodal densities (e.g., Figure 6b). Tarter and Lock's (1993, chap. 5) Fourier approach shows some promise with this problem, but further investigation is necessary to determine if it is worthwhile implementing in RTSYS.

It is hoped that future versions of RTSYS will implement maximum likelihood fitting of a range of three parameter distributions, including the displaced gamma, log-normal, inverse Gaussian, and Weibull. As reviewed by Heathcote (1994, see also Luce, 1986, and Ulrich & Miller, 1993), each distribution can be motivated by plausible psychological models, so selection of the best-fitting distribution may stimulate the investigators intuition about the mechanisms underlying RT in their paradigm. Significance of differences in fit can be tested with likelihood ratio tests. A potential extension involves fitting four parameter models, such as the convolution of normal and gamma distributions (Link, 1968) or the convolution of normal and inverse Gaussian distributions (Heath & Willcox, 1990). The convolved normal component may be useful in accounting for additive measurement noise extrinsic to the stochastic dynamics of the psychological process of interest (Heathcote, 1995a). Likelihood ratio test can be used to determine if the extra parameter significantly improves fits.

Fitting of parametric distributions may be improved by using the method of scoring. The method of scoring replaces the raw matrix of second partial derivatives of log likelihood used in fitting algorithms, such as Newton-

Raphson (e.g., Dobson, 1983, p. 30) and Marquardt-Levenberg, with its expected value under the model. Further investigation is required to determine whether the method of scoring speeds fitting, especially with more sophisticated search algorithms such as Marquardt-Levenberg. An added advantage of the method of scoring is that the asymptotic variance-covariance matrix of the parameters is equal to minus the expected second derivative matrix at the maximum likelihood estimate (also called the *information* matrix). Hence, the method of scoring provides estimates of parameter standard errors that allow inferential comparison of individual subject and Vincentized distribution parameters.

One criticism of likelihood ratio tests and standard error estimates given by the information matrix is that they rely on large sample results. With smaller samples, they may, therefore, be misleading. A recently developed solution is to use resampling and jackknife procedures to estimate standard errors and confidence intervals (e.g., Efron & Tibshirani, 1993). It is hoped that future versions of RTSYS will implement automatic estimation of these quantities. As described in the section on Vincentizing, the present version can be used for resampling with some preparation of data sets external to RTSYS. Given the computational demands of this approach, efficient fitting becomes an even more pressing issue, reinforcing the need to investigate and implement approaches such as the method of scoring. A parametric study using real data and simulated data to investigate the difference between resampled and asymptotic standard error estimates would also be of interest.

To facilitate the reader's understanding of the ex-Gaussian distribution, the Appendix contains Mathematica code for performing the convolution (Equation 4), plotting pdfs (Figure 2), sampling from the ex-Gaussian, and calculating and plotting likelihood surfaces (Figure 3). Similar graphics can be generated for other distributions defined in the Mathematica *ContinuousDistributions* package with minor alterations to the Mathematica code supplied.

Copies of RTSYS, including a hardcopy of the manual, can be obtained by sending an international money order for \$25 (US) (note that checks drawn on non-Australian banks are expensive to cash; if you wish to send such a check, please send \$30 [US] or \$20 (Australian) for orders within Australia. While many of the improvements suggested above will be implemented in pursuit of research issues, distribution of documented, robust, and user-friendly versions is time consuming and depends on the interest of members of the psychological research community. Please send bug reports, suggested improvements, and information on the use of RTSYS in research and teaching projects to rtsys@baal.newcastle.edu.au.

REFERENCES

- Box, G. E. P., & Cox, D. R. (1964). An analysis of transformations (with discussion). *Journal of the Royal Statistical Society B*, *26*, 211-246.
- Box, G. E. P., & Cox, D. R. (1982). An analysis of transformations re-
- visited, rebutted. *Journal of the American Statistical Association*, *77*, 209-210.
- BURBECK, S. L., & LUCE, R. D. (1982). Evidence from auditory simple reaction times for both change and level detectors. *Perception & Psychophysics*, *32*, 117-133.
- COX, D. R., & HINKLEY, D. V. (1974). *Theoretical statistics*. London: Chapman & Hall.
- D'ANGOSTINO, R. B., & STEPHENS, M. A. (1986). *Goodness-of-fit techniques*. New York: Dekker.
- DAWSON, M. R. W. (1988). Fitting the ex-Gaussian equation to reaction time distributions. *Behavior Research Methods, Instruments, & Computers*, *20*, 54-57.
- DOBSON, A. J. (1983). *An introduction to statistical modelling*. London: Chapman & Hall.
- EFRON, B., & TIBSHIRANI, R. J. (1993). *An introduction to the bootstrap*. New York: Chapman & Hall.
- HEATH, R. A., & WILLCOX, C. H. (1990). A stochastic model for interkeypress times in a typing task. *Acta Psychologica*, *75*, 13-39.
- HEATHCOTE, A. (1988). Screen control and timing routines for the IBM microcomputer family using a high-level language. *Behavior Research Methods, Instruments, & Computers*, *20*, 289-297.
- HEATHCOTE, A. (1994, September). *Reaction time distribution*. Seminar presented to the Department of Psychology, University of New South Wales, Sydney, Australia.
- HEATHCOTE, A. (1995a). *Measurement noise and models of cognition*. Paper presented at the Third Australian Cognitive Science Conference, Brisbane, Australia.
- HEATHCOTE, A. (1995b). A neural model of speed and accuracy. In *Proceedings of the Sixth Australian Conference on Neural Networks* (pp. 237-240). Sydney, Australia: University of Sydney, Department of Electrical Engineering.
- HEATHCOTE, A., POPIEL, S. J., & MEWHORT, D. J. K. (1991). Analysis of response time distributions: An example using the Stroop task. *Psychological Bulletin*, *109*, 340-347.
- HILGARD, E. (1938). A summary and evaluation of alternative procedures for the construction of Vincent curves. *Psychological Bulletin*, *35*, 282-297.
- HOCKLEY, W. E. (1984). Analysis of response time distributions in the study of cognitive processes. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *6*, 598-615.
- HOHLE, R. H. (1965). Inferred components of reaction times as functions of foreperiods duration. *Journal of Experimental Psychology*, *69*, 382-386.
- JOHNSON, D. M. (1939). Confidence and speed in the two-category judgment. *Archives of Psychology*, *241*, 1-53.
- KENDALL, M. G., & BUCKLAND, W. R. (1967). *A dictionary of statistical terms*. New York: Hafner.
- KENDALL, M. G., & STUART, A. (1967). *The advanced theory of statistics: Vol. 2. Inference and relationship* (2nd ed.). London: Griffin.
- KENDALL, M. G., STUART, A., & ORD, J. K. (1987). *Kendall's advanced theory of statistics: Vol. 1. Distribution theory* (5th ed.). New York: Oxford University Press.
- KENNEDY, W. J., JR., & GENTLE, J. E. (1980). *Statistical computing*. New York: Dekker.
- LINK, S. W. (1968). *Bisensory choice reactions*. Unpublished doctoral dissertation, Stanford University.
- LINK, S. W. (1982). Correcting response measures for guessing and partial information. *Psychological Bulletin*, *92*, 469-486.
- LUCE, R. D. (1986). *Response times: Their role in inferring elementary mental organization*. New York: Oxford University Press.
- MCGILL, W. J. (1963). Stochastic latency mechanisms. In R. D. Luce & R. R. Bush (Eds.), *Handbook of mathematical psychology* (Vol. 1, pp. 309-360). New York: Wiley.
- MEIJERS, L. M. M., & EIJKMAN, E. G. J. (1974). The motor system in simple reaction time experiments. *Acta Psychologica*, *38*, 367-377.
- MEWHORT, D. J. K., BRAUN, J. G., & HEATHCOTE, A. (1992). Response time distributions and the Stroop task: A test of the Cohen, Dunbar, and McClelland (1990) model. *Journal of Experimental Psychology: Human Perception & Performance*, *18*, 872-882.
- OLLMAN, R. (1966). Fast guesses in choice reaction time. *Psychonomic Science*, *6*, 155-156.

- 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). Hillsdale, NJ: Erlbaum.
- PRESS, W. H., FLANNERY, B. P., TEUKOLSKY, S. A., & VETTERLING, W. T. (1988). *Numerical recipes: The art of scientific computing*. Cambridge: Cambridge University Press.
- RATCLIFF, R. (1978). A theory of memory retrieval. *Psychological Review*, **85**, 59-108.
- RATCLIFF, R. (1979). Group reaction time distributions and an analysis of distribution statistics. *Psychological Bulletin*, **86**, 446-461.
- RATCLIFF, R. (1993). Methods for dealing with reaction time outliers. *Psychological Review*, **114**, 510-532.
- RATCLIFF, R., & MURDOCK, B. B. (1976). Retrieval processes in recognition memory. *Psychological Review*, **83**, 190-214.
- SILVERMAN, B. W. (1986). *Density estimation for statistics and data analysis*. London: Chapman & Hall.
- SMITH, P. (1995). Psychophysically-principled models of visual simple reaction time. *Psychological Review*, **102**, 567-593.
- TARTER, M. E., & LOCK, M. D. (1993). *Model-free curve estimation*. New York: Chapman & Hall.
- TUKEY, J. W. (1960). A survey of sampling from contaminated distributions (pp. 448-485). In I. Olkin (Ed.), *Contributions to probability and statistics*. Stanford, CA: Stanford University Press.
- ULRICH, R., & MILLER, J. (1993). Information processing models generating Lognormally distributed reaction times. *Journal of Mathematical Psychology*, **37**, 513-525.
- ULRICH, R., & MILLER, J. (1994). Effects of outlier exclusion on reaction time analysis. *Journal of Experimental Psychology: General*, **123**, 34-80.
- VAN ZANDT, T., & RATCLIFF, R. (1993, August). *Mimicking multiple processes with single-process models: Mixtures and parameter drift*. Paper presented at the 26th Annual Mathematical Psychology Meeting, Norman, OK.
- VINCENT, S. B. (1912). The function of the vibrissae in the behaviour of the white rat. *Behaviour Monographs*, **1** (Whole No. 5).
- YELLOT, J. I., JR. (1967). Correction for guessing in choice reaction time. *Psychonomic Science*, **8**, 321-322.

NOTES

1. When fitting fails, the barrier will usually cause search to converge to a value near it (i.e., σ or $\tau = \text{mean}/1,000$). RTSYS interprets any converged value of σ or τ less than twice the barrier ($\text{mean}/500$) as indicating a failure and sets all ex-Gaussian parameters to the missing value. Not all failed calls to the log likelihood function occur due to zero or negative value of σ or τ ; they may also occur for evaluations of the $\Phi(x)$ approximation when $|x| > 7$, or numerical evaluation of the integral on the left of Equation 5 is ≤ 0 or the integral on the right of Equation 5 is ≥ 1 . When such failures occur, minus log likelihood is also set to a very large barrier value. Fitting, however, is continued, because excursions into these regions may occur on the way to an adequate fit. Numerical problems are reported in both interactive and normal fitting modes so that suspect fits containing many numerical problems may be identified. When performing long fitting runs, the user is advised to echo fitting feedback to the printer so that suspect fits containing many numerical evaluation problems may be identified.

2. Unfortunately, the integrals in Equation 5 do not have closed form solutions or approximations and, hence, must be evaluated with computationally expensive numerical integration. Fortunately, the integration has to be performed only twice for each evaluation of L since the result does not vary with the value of each RT. The integration is performed using Press et al's. (1988, p. 116-118) Romberg integration on an open interval from 0 to a or b with an absolute accuracy of 1 in 10^6 . A lower limit of 0 is used because computation is faster for a definite integral and little probability density occurs for $RT < 0$ in RT data. Integration is performed using the midpnt function, and JMAX is set to 10 rather than the recommended 15 to save time on slowly converging integrals that occur when estimates of σ and τ are very small. If integration fails (JMAX is exceeded), $-\ln(L)$ is set to a large value.

APPENDIX

This appendix describes code used to perform the ex-Gaussian convolution and produce Figures 2 and 3 using Mathematica Version 2.2.1 for Microsoft Windows on a 25-MHz 486SL. To access predefined functions for statistical distributions, you must load the standard Mathematica package *ContinuousDistributions*:

```
<<Statistics'ContinuousDistributions'
```

Using this code as a template, you can explore the properties of the ex-Gaussian distribution (and other standard distributions defined in *ContinuousDistributions*).

To define the ex-Gaussian pdf, you can either type in the function given in the text,

```
exgaus[x_,m_,s_,t_] := Exp[(s^2/(2*t^2))-
  ((x-m)/t)]*(t*(2*Pi)^0.5)^-1*
  Integrate[Exp[-y^2/2],
    {y,-Infinity,((x-m)/s)-(s/t)}],
```

or directly convolve the component normal and exponential distributions defined in the package *ContinuousDistributions*,

```
pevg[m_,s_,t_,x_] :=
  Integrate[PDF[ExponentialDistribution[(1/t)],z]*
  PDF[NormalDistribution[m,s],(x-z)],{z,0,Infinity}].
```

Note that predefined exponential distribution parameter, b , is the inverse of τ . Hence, the parameter t of the function *pevg* is entered in the exponential density as $1/t$. The integration is performed over values where the exponential has nonzero probability density ($0, \infty$). The integral in the output of both versions of the ex-Gaussian density is expressed in terms of the error function (Erf), which can be efficiently approximated numerically.

Figure 2 was produced by plotting the predefined normal and exponential densities and the ex-Gaussian density defined above:

```
Plot[PDF[NormalDistribution[500,50],x],{x,300,700},
  PlotRange->{0,0.01},PlotStyle->AbsoluteThickness[1]]
Plot[PDF[ExponentialDistribution[0.01],x],{x,1,500},
  PlotRange->{0,0.01},PlotStyle->AbsoluteThickness[1]]
Plot[pevg[500,50,100,x],{x,300,1000},
  PlotStyle->AbsoluteThickness[1]].
```

Samples from the ex-Gaussian were obtained as the sum of samples from the normal and exponential distributions:

```
randexg[m_,s_,t_] := Random[NormalDistribution[m,s]] +
  Random[ExponentialDistribution[1/t]].
```

The sample of 10 values was obtained with the following commands (note that this command will produce different values each time it is invoked):

```
sample = Table[randexg[500,50,100],{i,10}].
```

The ex-Gaussian likelihood function was obtained by summing minus the natural log of the ex-Gaussian density over the set of 10 sample values:

```
ll[m_,s_,t_,rtset_] := Apply[Plus,Log[pevg[m,s,t,rtset]]].
```

Plots of the likelihood function in the neighborhood of the sample generating function's parameters were obtained with the following commands. Note that the plots were very computationally intensive (e.g., several hours each for the three-dimensional plots), so the two-dimensional plots were saved and displayed

using the Show function to allow selection of the appropriate plot regions without recomputation. Given the shape of the two-dimensional plot regions, the three-dimensional plots were computed directly.

```

plotlmtmu = Plot[lI[m,50,100,sample],{m,460,520}]
Show[plotlmtmu,PlotRange -> {{460,520},{57,58}},
     AxesOrigin -> {460,57}]
plotltsig = Plot[lI[500,s,100,sample],{s,10,80}]
Show[plotltsig,PlotRange -> {{10,80},{56,59}},
     AxesOrigin -> {10,56}]

```

```

plotlltau = Plot[lI[500,50,t,sample],{t,40,120}]
Show[plotlltau,PlotRange -> {{40,120},{56.5,57.8}},
     AxesOrigin -> {40,56.5}]
Plot3D[lI[500,s,t,sample],{s,10,80},{t,40,120}]
Plot3D[lI[m,50,t,sample],{t,40,120},{m,450,550}]
Plot3D[lI[m,s,100,sample],{s,10,80},{m,450,550}]

```

(Manuscript received January 3, 1994;
revision accepted for publication March 8, 1995.)