2024 ■ Vol. 20 ■ no. 2

Response time distribution analysis of medium-sized datasets in MATLAB

Jarrod Blinch [a] and Denis Cousineau [b]

[a]Department of Kinesiology & Sport Management, Texas Tech University
[b]School of Psychology, University of Ottawa

Abstract ■ Response time data have a positively skewed distribution. The challenge with this is that a measure of central tendency and dispersion does not adequately describe a skewed distribution. A researcher relying on only response time mean and standard deviation could make incorrect conclusions about response time. The best way to analyze response time data is with a distribution analysis. One reason that response time distribution analyses are atypical is that at least 100 trials are recommended per participant and condition. In the current tutorial, we demonstrate a distribution analysis technique that requires as few as 40 participants with 40 trials per condition. This technique involves geometric quantile averaging (GQA) and the quantile maximum probability estimator (QMPE). Each step of the analysis is detailed with a MATLAB script, flexible MATLAB functions, and experimental response time data. Our goal was to lower the barriers to entry for response time distribution analysis so that more researchers will choose to thoroughly examine response time data.

Keywords ■ ex-Gaussian distribution function, geometric quantile averaging (GQA), quantile maximum probability estimator (QMPE). **Tools** ■ MATLAB.

jarrod.blinch@gmail.com

10.20982/tqmp.20.2.p106

Acting Editor ■ Roland Pfister (Universität Würzburg)

Reviewers
■ Josh Sandry (Monclair State University)
■ and one anonymous reviewer.

Introduction

How should we analyze response time, the time from the go signal to movement initiation?[1] Traditional response time analysis involves a measure of central tendency (e.g., mean or median; c.f. Miller, 1988) and a measure of dispersion (e.g., standard deviation or variance). This type of analysis is appropriate for variables that have a Gaussian (normal) distribution function. Response time data, however, almost always have a positively skewed distribution function. For such data, the traditional analysis is better replaced by a distribution analysis. One potential reason that response time distribution analyses are atypical is that at least 100 trials are recommended for each participant and in each condition (per cell; Lacouture & Cousineau, 2008; Ratcliff, 1979; Van Zandt, 2000), and most studies have far fewer trials than that. In the current tutorial, we demonstrate a distribution analysis technique that requires fewer trials (as few as 40 participants with 40 trials per cell). This technique involves two steps: calculating group quantiles by merging individual quantiles and then calculating the group distribution by fitting the ex-Gaussian distribution function to the group quantiles. The current best techniques for these two steps are geometric quantile averaging (GQA; Cousineau et al., 2016) and the quantile maximum probability estimator (QMPE; Heathcote et al., 2004), respectively. We will return to these steps after further justification for response time distribution analysis and the use of the ex-Gaussian distribution function.

The challenge with response time data is that a measure of central tendency and dispersion does not adequately describe a skewed distribution function. In the worst case, two positively skewed distribution functions can be visibly different and yet have identical means and standard deviations. Figure 1, left, shows two visibly different distribution functions with identical means ($\bar{X} = 1000$) and standard deviations ($s = 100.0$). A researcher relying on only means and standard deviations would incorrectly conclude that

[1]Response time and reaction time have both been used to describe the duration of time from the go signal to movement initiation (Luce, 1986; Welford, 1980). We decided to use response time in the current tutorial because it is the more commonly used term in the methodology of distribution analysis.

Figure 1 ■ Example ex-Gaussian Distribution Functions. (Left) The two ex-Gaussian distribution functions are visibly different and yet their means (1,000) and standard deviations (100.0) are identical. This example demonstrates the need to analyze response time distributions and not just the means and standard deviations. The two distribution functions do have different skew (0.71 grey line and 1.71 black line), but stable estimates of skew are impractical in most experiments as they require hundreds of trials. (Right) Three examples of ex-Gaussian distribution functions. The solid grey line is a typical ex-Gaussian distribution function with positive skew. The black line is an ex-Gaussian distribution function with a τ parameter of zero, which is identical to a Gaussian (normal) distribution function. The dotted grey line is an ex-Gaussian distribution function with a σ parameter of zero, which is identical to an exponential distribution function.

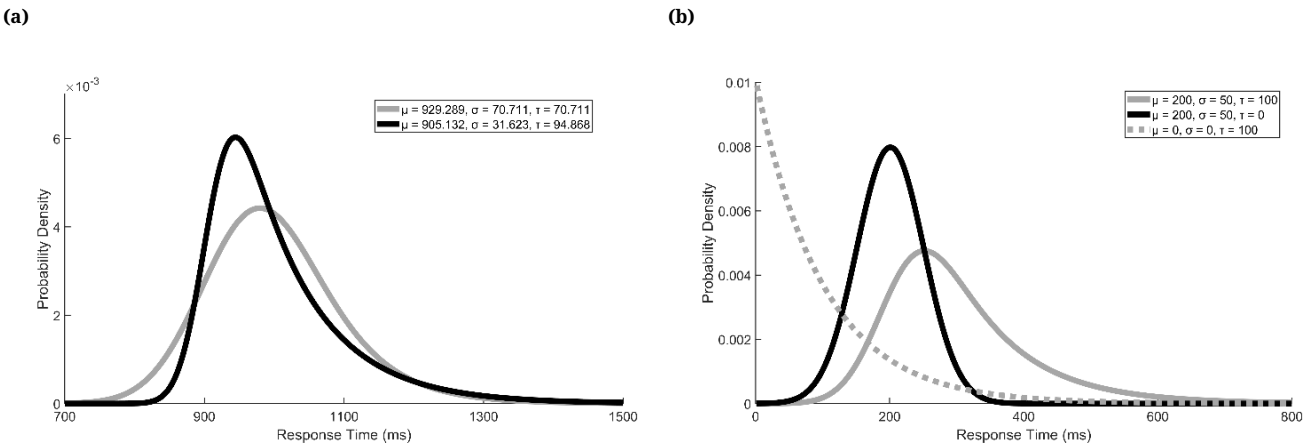

response time of the two conditions is not significantly different. A skewed distribution function can be described by a measure of central tendency, dispersion, and skewness. However, there are two major issues with measurements of skewness (Heathcote et al., 1991; Ratcliff, 1979). First, the sampling variance associated with skewness is extremely large. Consequently, a stable estimate of skewness requires hundreds of trials per cell. Second, measurements of skewness are sensitive to outliers even with many trials. The best way to analyze response time data is, therefore, with a distribution analysis (Whelan, 2008).

The first decision in a distribution analysis is choosing a theoretical distribution. The Log-normal, Weibull, and Wald distributions have occasionally been used to fit positively skewed response time data. The most popular choice, however, is the exponentially modified Gaussian (ex-Gaussian) distribution function, which was popularized in response time research by Hohle (1965), Ratcliff (1978, 1979), and Luce (1986). The ex-Gaussian distribution function is derived by the combination of a Gaussian distribution function and an exponential distribution function (in mathematical terms, it is the convolution of those two functions). An ex-Gaussian distribution function is described by three parameters: mu (μ), the mean of the normal component, sigma (σ), the standard deviation of the normal component, and tau (τ), a single value for the mean and standard deviation of the exponential component. It is the τ parameter that defines the positive skew of the ex-Gaussian distribution. The mean and standard deviation of an ex-Gaussian distribution are $\mu + \tau$ and $\sqrt{\sigma^2 + \tau^2}$, respectively. These formulas demonstrate the results of combining the Gaussian and exponential distributions. Three examples of ex-Gaussian distribution functions are shown in Figure 1, right.

The μ, σ, and τ parameters are typically used to describe the distribution of the response time data and to then make comparisons between conditions, which we demonstrate in the current tutorial. An advantage of this descriptive interpretation of the ex-Gaussian parameters is that it can be applied broadly to all sorts of response time experiments. One such example, using older methodology, is Heathcote et al.'s (1991) application of the ex-Gaussian distribution function to response time in the Stroop task. The generalizability of the descriptive approach comes with a disadvantage; it lacks a theoretical interpretation of the ex-Gaussian parameters. There are some studies that connect ex-Gaussian parameters to cognitive processes (e.g., Hockley, 1984; Ratcliff, 1978; Schmiedek et al., 2007), but they do not generalize beyond the specific response time paradigm investigated.

Regardless of whether a descriptive or theoretical interpretation is used, the ex-Gaussian distribution function must first be fit to the response time data. If you are in the fortunate situation of having at least 100 trials per cell, then you can fit an ex-Gaussian distribution function to each cell. There is a tutorial on this technique for MATLAB by Lacou-

ture and Cousineau (2008) and for Python by Moret-Tatay et al. (2018). This technique will yield an estimate of μ, σ, and τ for each participant and condition, which can then be statistically analyzed with, for example, an analysis of variance on each parameter. One cannot reliably fit the ex-Gaussian distribution function to individual distributions with fewer than 100 trials (Lacouture & Cousineau, 2008; Ratcliff, 1979; Van Zandt, 2000); however, one can fit the distribution function to the group distribution (Cousineau et al., 2016). A group distribution, in this case, is when the distribution of each participant (in a given condition) is merged into an aggregate "group" distribution. The basic procedure involves two steps: calculating group quantiles by merging the individual quantiles for each cell (each participant and condition) and then calculating the group distribution by fitting the ex-Gaussian distribution function to the group quantiles.

A popular technique for calculating a group distribution was introduced by Ratcliff (1979). Ratcliff calculated group quantiles by taking the arithmetic mean of the individual quantiles. The problem with arithmetic quantile averaging (Vincent averaging or Vincentizing) is that it can introduce artifacts in the group distribution that are not present in any of the individual distributions even with 100 trials per cell (Rouder & Speckman, 2004; Van Zandt, 2000). The current best technique to estimate group quantiles is GQA. It preserves the characteristics of the individual quantiles without introducing artifacts (see Cousineau et al., 2016, for the development of GQA). GQA used in the current tutorial has the added benefit that 95% confidence intervals for the group quantiles can be computed. We show how to calculate adjusted confidence intervals that allow for visual analysis. This relies on the golden rule of adjusted confidence intervals, which states that if a mean is outside the 95% adjusted confidence interval of another mean, then the means are likely statistically significant at the .05 level (see Cousineau, 2017; Cousineau et al., 2021, for details on adjusted confidence intervals). Visual analysis is especially important when comparing group distributions as we will show that statistical analysis cannot be used.

With the group quantiles calculated with GQA, the ex-Gaussian distribution function can then be fit to the group quantiles with QMPE. QMPE returns unbiased parameters with low variability with as few as 40 trials per cell (Brown & Heathcote, 2003; see Heathcote et al., 2002, 2004, for the development of QMPE). Its main advantage is that it is uninfluenced by response time outliers. In the current tutorial, we will demonstrate GQA and QMPE in MATLAB. Our goal was to lower the barriers to entry for response time distribution analysis so that more researchers will choose to thoroughly examine response time data.

Method

The data for the present analysis were previously reported by Blinch et al. (in press, Experiment 1). The relevant details from that study are summarized below.

Participants, Design, Apparatus, Procedures, and Data Analysis

Forty-one volunteer participants (27 female and 14 male participants) were recruited from the Texas Tech University community. The age of participants ranged from 19 to 37 years old (M = 22.0, Mdn = 21.0, SD = 3.35). Participants completed a two-choice response time task with pointing movements to either a short- or long-distance target (10 or 20 cm, respectively) in two conditions. In one condition, visual information was available for the entire trial (vision condition). In the other condition, visual information was occluded during movement execution (no vision condition). Each condition consisted of 128 pseudorandomized trials, with 64 trials to the short-distance target and 64 trials to the long-distance target.

Participants were seated at a table with a button box on the surface of the table. The button box consisted of a home button, a short-distance target, and a long-distance target. The home button was a micro push-button switch. The short- and long-distance targets were clear push-button switches. The push-button switches were each illuminated by a diffused green light-emitting diode. Participants wore visual occlusion spectacles (Translucent Technologies, PLATO) that controlled their access to visual information during the trials.

All trials began with the spectacles open to allow visual information. The participant started each trial by pressing and holding down the home button with the index finger of their dominant hand. There was a 1-2 s variable foreperiod before either the short- or long-distance target illuminated as the go signal. The participant was instructed to "react and press the illuminated target as quickly and as accurately as possible". The participant held down the target button at the end of their movement until the 2-s recording interval ended and the target darkened. They could then begin the next trial by returning to the home button. Trials in the condition without visual information during movement execution were slightly different, in that the spectacles closed to occlude visual information from the release of the home button until the target button was pressed (i.e., from movement initiation to movement termination). The spectacles opened at movement termination to provide terminal feedback of the movement. Response time on each trail was calculated as the interval of time from illumination of the target button to release of the home button, that is, from the go signal to movement initiation.

Listing 1 ■ GQA Steps 1 to 5 From analyse.m

```
%% Geometric quantile averaging (GQA) Step 1
% Calculate base response time for each participant and condition.
base_rt_array = get_base_rt(rt_array, 1);

%% GQA Step 2
% Calculate quantiles for each participant and condition.
quantile_total = 20;
quantile_array = get_quantiles(rt_array, quantile_total, 1);
quantile_unmodified_array = quantile_array;

%% GQA Step 3
% Subtract base RT from the quantiles for each participant and condition.
for participant_num = 1:participant_total
    quantile_array(:,:,participant_num) = quantile_array(:,:,participant_num) -
   base_rt_array(participant_num,:);
end

%% GQA Step 4
% Calculate the geometric mean of each quantile across participants for each condition.
quantile_geometric_mean_array = geomean(quantile_array, 3);

%% GQA Step 5
% Add the mean base RT for each condition.
quantile_geometric_mean_array = quantile_geometric_mean_array + mean(base_rt_array,
    1);
```

Results

In this section, we detail the process of response time distribution analysis with GQA and QMPE with the MATLAB script analyse.m. The first section of code in analyse.m (denoted by the two comment characters [%%]) cleans up the MATLAB environment (closing all figure windows, removing all variables, functions, MEX links, and clearing the command window). The second code section loads the response time data from the `rt_array.mat` file into the `rt_array` variable, which is a 3-D matrix with 64 rows, 4 columns, and 41 pages. Each page contains the response time data of one participant. The columns contain the four different conditions: vision during movement execution and a short-distance target (vision short), vision long, no vision short, and no vision long. Finally, the rows are the 64 trials.

GQA

GQA involves six steps. 1) Estimate base response time for each participant and condition; 2) Calculate quantiles for each participant and condition; 3) Subtract base response time from the quantiles; 4) Calculate the geometric mean of each quantile across participants for each condition; 5) Add the arithmetic mean base response time to each condition; 6) Calculate confidence intervals for each quantile in each condition.

In the GQA Step 1 code section (Listing 1), base response time is estimated for each participant and condition. Base response time is the shortest possible response when performing the task correctly. The true base response time is unknown and so it is estimated with the following equations (Cousineau et al., 2016),

$$Base\ RT = \frac{X_{1:n} - h\,(n,1)\,\bar{X}}{1 - h\,(n,1)} \tag{1}$$

In which

$$h\,(n,1) = 3 \times 2^{2n-1} \frac{(n!)^2}{(2n+1)!} \tag{2}$$

In Equations 1 and 2, n is the number of trials, $\bar{X}$ is the arithmetic mean, and $X_{1:n}$ is the short-

Table 1 ■ Output From Geometric Quantile Averaging

Step	2	3	4	5	6
Participant	10	10	All	All	All
Condition			No Vision Long (4)		
Variable	`quantile_array(:,4,10)`	`quantile_array(:,4,10)`	`quantile_geometric_mean_array(:,4)`	`quantile_geometric_mean_array(:,4)`	`quantile_geometric_95_ci_size_array(:,4)`
	250.2857	21.8347	29.9526	232.2144	4.9313
	259.7857	31.3347	37.6837	239.9454	5.6230
	265.9286	37.4776	42.4988	244.7606	6.0605
	…	…	…	…	…
	340.7857	112.3347	128.1106	330.3723	8.4243
	346.4048	117.9537	147.4293	349.6911	10.9943
	358.5238	130.0728	182.2889	384.5507	16.3093

Note. Data from only participant 10 and condition 4 are listed for Steps 2 and 3. Data from only condition 4 are listed for Steps 4, 5, and 6. Four decimal places are shown to match the output from MATLAB.

est response time. These equations were combined in MATLAB as `function [ output_array ] = get_base_rt(data_array, dimension)`. In analyse.m, the code for GQA Step 1 is `base_rt_array = get_base_rt(rt_array, 1);`. The variable `base_rt_array` is a 2-D matrix with 41 rows, one for each participant, and 4 columns, one for each condition; for example, base response time for the tenth participant in the no vision long condition, `base_rt_array(10,4)`, is 228.4510 ms. More information on the `get_base_rt` function, and the other functions, is available in MATLAB with the help command; for example, `help get_base_rt`.

In GQA Step 2, 20 quantiles (.0476, .0952, ..., .9048, .9524) are calculated for each participant and condition. We will explain how to determine the number of quantiles for any dataset in the Discussion. The quantiles are calculated with `function [ output_array ] = get_quantiles(data_array, number_of_quantiles, dimension)`. The main code in this section is `quantile_array = get_quantiles(rt_array, quantile_total, 1);`. `quantile_total` is 20 and `quantile_array` is a 3-D matrix with 20 rows (one for each quantile), 4 columns (one for each condition), and 41 sheets (one for each participant). Example output from this step, and the following GQA steps, is shown in Table 1.

Built-in MATLAB functions and operators are used for QGA Steps 3 to 6. In GQA Step 3, the base response time for each participant and condition is subtracted from the matching quantiles with a for loop. In GQA Step 4, the geometric means of the quantiles are calculated (`quantile_geometric_mean_array = geomean(quantile_array, 3);`). These are stored in `quantile_geometric_mean_array`, which is a 2-D matrix with 20 rows (one for each quantile) and 4 columns (one for each condition). In GQA Step 5, the group base response time in each condition is calculated as the arithmetic mean base response time across participants. These group base response times are then added to the geometric mean quantiles (`quantile_geometric_mean_array = quantile_geometric_mean_array + mean(base_rt_array, 1);`).

In GQA Step 6, 95% confidence intervals are calculated for the group quantiles in each condition and then the results are plotted. We recommend difference-adjusted confidence intervals for between-participant comparisons and correlation- and difference-adjusted confidence intervals for within-participant comparisons (Cousineau, 2017; Cousineau et al., 2021). More specifically, the difference adjustment accounts for comparing two confidence intervals and the correlation adjustment accounts for within-participant correlation across conditions for each quantile (but not across quantiles). With these adjusted confidence intervals, if a specific group quantile in one condition falls outside the 95% adjusted confidence interval of the same quantile in another condition, then they are likely significantly different at the .05 level (the golden rule of adjusted confidence intervals; Cousineau, 2017; Cousineau et al., 2021). Conversely, if a specific group quantile in one condition is within the 95% adjusted confidence interval of the same quantile in another condition, then they are likely not significantly different. Correlation- and difference-adjusted confidence intervals were used in the current tutorial to compare the group quantiles across the four conditions: vision short, vision long, no vision short, and no vision long. These confidence intervals were calculated with the Cousineau-Morey method. This method assumes that the means are normally distributed and that there are equal variances of difference scores. If there are issues with normality, then the visual analysis with the golden rule and the statistical analysis are more likely to differ. If there are issues with normality, then the confidence in-

Figure 2 ■ Group Quantiles in the Four Conditions. Error bars are 95% correlation- and difference-adjusted confidence intervals (Cousineau, 2017; Cousineau et al., 2021).

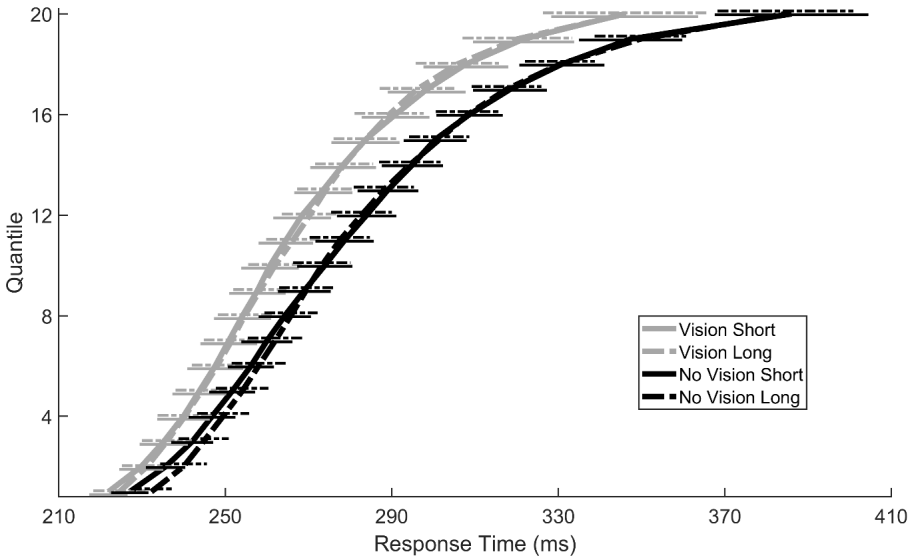

tervals likely have different widths. It may be necessary to calculate an average width for the confidence intervals (square the widths, calculate the arithmetic mean, and then take the square root) or use the Tryon adjustment (Tryon, 2001).

The first for loop in GQA Step 6 (Listing 2) removes between-participant variance from the quantiles, which is necessary for correlation-adjusted confidence intervals. The correlation- and difference-adjusted 95% confidence intervals are then calculated. Note that the standard error of the geometric mean is calculated with the following formula (Harding et al., 2014):

$$SE_G = G_x \times \frac{s_{\log x}}{\sqrt{n-1}} \tag{3}$$

In Equation 3, G_x is the geometric mean for one quantile and condition, $s_{\log x}$ is the standard deviation of the logarithm of the individual data for the same quantile and condition, and n is the number of participants. The halfwidth of the correlation- and difference-adjusted confidence interval is then calculated with the following formula:

$$CI_{HW} = \sqrt{2} t_{\alpha,n-1} SE_G \sqrt{k/k-1} \tag{4}$$

Note that $t_{\alpha,n-1}$ is the critical value of the Student's t-distribution and k is the number of conditions.

The second for loop in GQA Step 6 plots each participant's quantiles (red lines) and the group quantiles (black line) with correlation- and difference-adjusted 95% confidence intervals with each condition on its own figure. Figure 2 illustrates the group quantiles with correlation- and difference-adjusted 95% confidence intervals in all four conditions. The conditions can be visually analyzed with the golden rule of adjusted confidence intervals to determine whether statistically significant differences between conditions at each group quantile are likely. Applying the golden rule to Figure 2, the vision short and vision long conditions are likely comparable at every quantile. Likewise, the no vision short and no vision long conditions are likely comparable at every quantile. Most importantly, the vision conditions are likely significantly different from the no vision conditions from the .1429 to the .9524 quantiles. (At the .0476 and .0952 quantiles, the vision conditions are significantly different from the no vision long condition.) The differences between the vision and no vision conditions can be further explored by comparing their ex-Gaussian distribution functions.

QMPE

Fitting the ex-Gaussian distribution function to the group quantiles in each condition involves three steps. 1) Add zero milliseconds and positive infinity milliseconds as the 0 and 1 quantiles; 2) Estimate the parameters of the ex-Gaussian distribution function for each group/condition with QMPE; 3) Optionally, calculate the ex-Gaussian probability density functions for each group/condition and plot the resulting best-fitting distributions. In the QMPE Step 1 code section (Listing 3 at the end), zero and positive infinity are added as the first and last group quantiles. This is a requirement of the `QMPE` func-

Listing 2 ■ First Half of GQA Step 6 From analyse.m

```
%% GQA Step 6
% Plot the individual and group quantiles.
% For the group quantiles, include correlation – and difference –adjusted 95% confidence intervals.

% Correct quantile_array for correlation –adjusted Cis
quantile_corrected_array = ones(size(quantile_array)) .* NaN;

for x = 1:size(quantile_array,1)
    temp_array = squeeze(quantile_array(x,:,:))';

    participant_mean_array = nanmean(temp_array,2);
    grand_mean = nanmean(temp_array(:));

    temp_corrected_array = bsxfun(@minus, temp_array, participant_mean_array) +
   grand_mean;
    for y = 1:size(temp_corrected_array, 1)
        quantile_corrected_array(x,:,y) = temp_corrected_array(y,:);
    end
end

% Calculate the halfwidth of 95% confidence intervals
quantile_geometric_se_array = geomean(quantile_array, 3) .* (std(log(
   quantile_corrected_array), 0, 3) / sqrt(participant_total - 1));
quantile_geometric_95_ci_size_array = sqrt(2) * tinv(.975, participant_total - 1) *
    quantile_geometric_se_array * sqrt(condition_total/(condition_total-1));
```

tion, which requires the first and last quantiles to be the extremes of the theoretical distribution (zero and positive infinity for the ex-Gaussian distribution). The geometric mean quantiles in the no vision long condition, `quantile_geometric_mean_2_array(:,4)`, are 0, 232.2144, 239.9454, ..., 349.6911, 384.5507, and `Inf`, which is how MATLAB represents positive infinity.

In QMPE Step 2, the parameters of the ex-Gaussian distribution function (μ, σ, and τ) are estimated with `function [ parameter_array, minimum_log_likelihood, return_code ] = QMPE(data_array, distribution, varargin)`. The input distribution is set to exGaussian for the ex-Gaussian distribution function.[2] The input `varargin` is an optional variable-length input argument list for advanced options. It was set to `'plotF', 0, 'startPoint', [mu_init, sigma_init, tau_init]`. A `plotF` of 0 does not plot the real time results of parameter estimation during the search. The input `startPoint` sets the initial parameters for the search to `mu_init`, `sigma_init`, and `tau_init`, which are reasonable starting search estimates for each condition. These values for the no vision long condition (condition 4) are 247.0700, 36.8010, and 43.2242, respectively. As for the output, `parameter_array` is a row vector with μ, σ, and τ. The output `minimum_log_likelihood` is related to the index of fit, and `return_code` equals 1 on success. The parameters in the four conditions (`parameter_condition_array`) are listed in Table 2. The QMPE function uses the Nelder-Mead method for multidimensional unconstrained optimization (MATLAB's `fminsearch` function).

In QMPE Step 3, the densities of the entire ex-Gaussian distribution function are determined with `function pdf = exGaussianpdf(x, mu, sigma, tau)`. The probability density functions are calculated for response time values from 0 to 1,000 ms (in increments of 1 ms; `0:1000` in MATLAB) with the `exGaussianpdf` function. These are stored in `pdf_array`, a 2-D matrix with 1,001 rows (one for each response time value) and 4 columns (one for each condition). Finally, the four ex-Gaussian probability density functions are plotted on the

[2] QMPE was originally written in Fortran 90 (Heathcote et al., 2004). It was then written in MATLAB by Valerio Biscione (www.mathworks.com/matlabcentral/fileexchange/46330-qmle-zip). We edited Valerio's code to use cumulative distribution functions for faster computation.

Table 2 ■ Parameters of the ex-Gaussian Distribution Functions of the Group Distribution in the Four Conditions

Condition	μ (ms)	σ (ms)	τ (ms)
Vision Short	235.6180	17.0159	35.0118
Vision Long	236.8007	16.4027	34.1507
No Vision Short	240.9808	19.3727	44.8679
No Vision Long	242.0632	15.7175	44.8294

Note. Note. Four decimal places are shown to match the output from MATLAB.

Figure 3 ■ Best-Fitting ex-Gaussian Probability Density Functions Based on the Group Quantiles in the Four Conditions

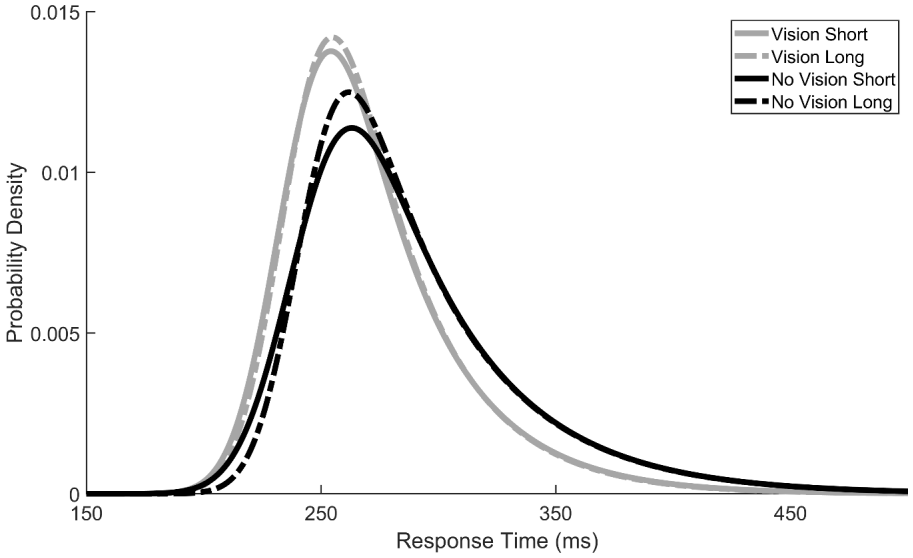

same figure (Figure 3). The excellent fit of the ex-Gaussian distribution functions to the group quantiles is depicted in Figure 4.

It is impossible to statistically compare the parameters of the ex-Gaussian distribution functions in the four conditions as these estimates have unknown standard errors. However, we know from the group quantiles and their adjusted confidence intervals that there are significant differences between the vision and no vision group distributions. These significant differences can aid in the visual analysis of the ex-Gaussian parameters. The mean of the Gaussian component (μ) in the no vision conditions (242 ms) was longer than in the vision conditions (236 ms). This is illustrated in the probability density functions as a rightward shift in the mode for the no vision conditions (Figure 3, black lines vs. grey lines). A larger μ suggests that response time on most trials is longer without vision during movement execution. The standard deviation of the Gaussian component (σ) was 16.7 ms in the vision conditions. In comparison, σ was larger for short movements in the no vision condition (19.4 ms). This is illustrated as a slightly longer left tail and a lower mode for the no vision short condition (Figure 3, solid black line). A larger σ suggests that response time is more variable for short movements without vision. Finally, the single parameter for the mean and standard deviation of the exponential component (τ) was larger in the no vision conditions (44.8 ms) compared to the vision conditions (34.6). This is illustrated as longer right tails and lower modes for the no vision conditions (Figure 3, black lines vs. grey lines). A larger τ suggests that trials with long response time are more common when there is no vision during movement execution. Overall, response time performance is worse without vision: it is longer, more variable for short-amplitude movements, and there are more trials with long response time. Note that in our comparisons of the parameters with and without vision, we focused on the single most obvious difference between the ex-Gaussian probability density functions. However, changing a single parameter will have obvious and subtle effects on the entire density function. A larger σ or τ, for example, will also decrease the peak of the probability density function, which is visible in Figure 3.

Figure 4 ■ Best-Fitting ex-Gaussian Cumulative Density Functions Compared to the Group Quantiles in the Four Conditions

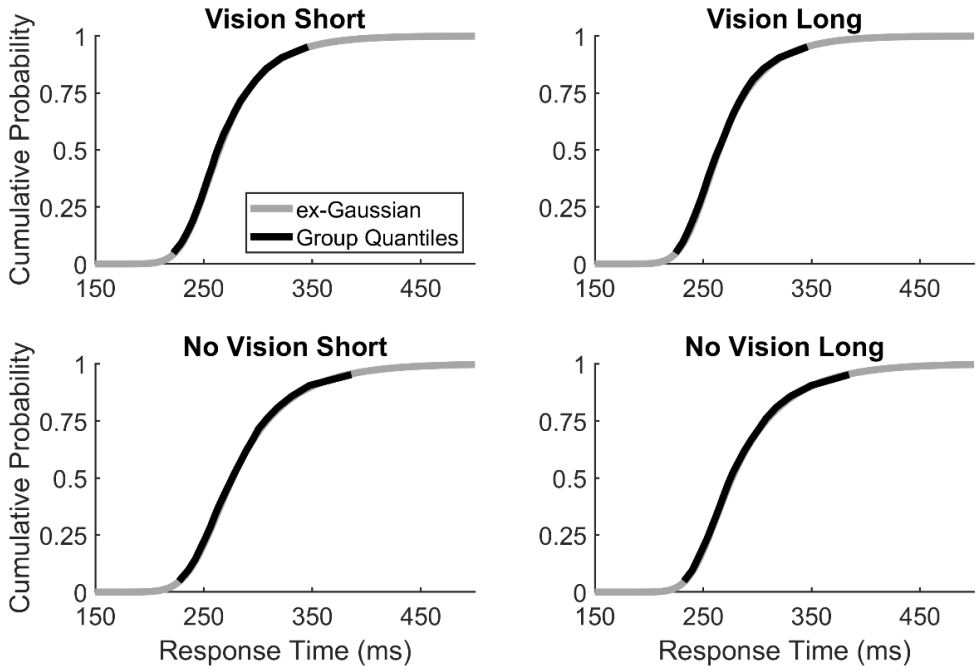

Discussion

In the current tutorial, we demonstrated a procedure for the distribution analysis of response time data. (Recall that a distribution analysis is necessary to properly analyze response time data because it is positively skewed.) The procedure involves calculating group quantiles with GQA and then fitting the ex-Gaussian distribution function to the group quantiles with QMPE. An advantage of this technique is that it requires as few as 40 trials per cell. More specifically, two studies investigated the bias and efficiency of estimating ex-Gaussian distribution functions with QMPE (Heathcote et al., 2002, 2004). Heathcote et al. examined estimates based on 40, 80, and 160 trials, and they found robust and efficient estimates even with 40 trials.

There are two important caveats when applying these studies to the current tutorial. First, bias and efficiency with fewer than 40 trials have not been examined. Second, and most importantly, these studies investigated the use of QMPE on individual participants. The present tutorial involves GQA followed by QMPE. To address the second caveat, we ran several simulations to estimate the efficiency of GQA followed by QMPE (detailed in the Appendix). The simulations suggested that the geometric mean of quantiles from 40 participants with 40 trials per condition was almost as efficient as maximum-likelihood estimation of a single participant with 100 trials (Van Zandt, 2000). Thus, GQA followed by QMPE on a moderate number of participants with a moderate number of trials has the potential to return efficient estimates.

For typical response time distribution analysis, where, for instance, the ex-Gaussian distribution function is fit to individual distributions (not the group distribution), 100 trials per cell are recommended (Lacouture & Cousineau, 2008; Ratcliff, 1979; Van Zandt, 2000). Lacouture and Cousineau (2008) showed how to fit an ex-Gaussian distribution to response time data. They examined the bias and efficiency of their method with a Monty Carlo study based on 10, 20, 30, 50, 100, 500, 1000, and 5000 trials. As seen in Figure 7 of Lacouture and Cousineau (2008), the magnitude of the bias of the parameter estimates doubles from a sample size of 500 to 100 and then doubles again from 100 to 50. The variance of the estimates was about 20 to 40 ms with a sample size of 100. This is well below the observed variability in response time, and so the added uncertainty of the estimates is acceptable. With a sample size of 50, the added variability might exceed the standard deviation of response time in some cells, which should be avoided. With a sample size of 500, the added uncertainty is small, but such large samples are prohibitive. Hence, a sample size of 100 is an acceptable tradeoff. Collecting 40 trials per cell is

certainly more feasible than 100, but 40 participants and 40 trials per cell could still be more than some studies provide. Thus, the number of trials required should be considered when designing an experiment that includes distribution analysis.

The number of trials will also affect the number of quantiles/bins. As with the number of trials, the more quantiles/bins, the better (Heathcote et al., 2002). Let us consider just the number of bins for simplicity's sake (the number of quantiles can be calculated by subtracting one from the number of bins). There were 64 trials in each cell in the present dataset. The number of bins is maximized when there is one datapoint in each bin. Therefore, the 64 trials per cell in the present dataset could be used to create at most 64 bins. However, that is assuming there are no outliers in the data, which is unlikely. When there are outliers, there needs to be more than one datapoint in each bin to reduce the impact of outliers. A possible formula to determine the number of bins based on the number of trials (n) and the hypothesized proportion of outliers (p_o) is as follows,

$$Bins = \frac{1}{p_o + \frac{1}{n}} \tag{5}$$

We estimated the proportion of outliers in each cell by applying a modified square root transformation and then using z-scores to identify outliers ($\alpha = .01$ and a Bonferroni correction based on the number of trials in each cell; Cousineau & Chartier, 2010). This procedure is demonstrated on the current dataset in the `proportion_of_outliers.m` MATLAB script. The results suggested that the proportion of outliers ranged from .00 to .03125. We used the bin with the most outliers ($n = 64, p_o = .03125$) to calculate the number of bins, 21.33, which was rounded down to the nearest whole number. Thus, the number of bins in the current tutorial was 21 and the number of quantiles was 20. Two hypothetical examples are when the proportion of outliers is 0 or .4844, the number of bins is 64 or 2 (the 2 bins are cut by the .5 quantile: the median).

There are three potential downsides to group distributions compared to individual distributions (Lacouture & Cousineau, 2008). First, more calculations are required for a group distribution. However, we have encapsulated these calculations into MATLAB code that is publicly available at osf.io/be7gp/ . Second, group distributions cannot be analyzed with, for example, analysis of variance or regression. With the group distribution, one can estimate the parameters of the ex-Gaussian distribution function. However, it is impossible to estimate these parameters for each participant with the group distribution. Recall that at least 100 trials per cell are recommended to fit individual distributions, which could then be analyzed with parametric statistics. Our suggested solution is to use the golden rule of confidence intervals to compare different group distributions to estimate statistical significance (Cousineau, 2017; Cousineau et al., 2021). For the golden rule, correlation- and difference-adjusted confidence intervals must be used for within-participant comparisons and difference-adjusted confidence intervals must be used for between-participant comparisons. The present tutorial demonstrated how to interpret correlation- and difference-adjusted confidence intervals in a within-participant design. The third downside to group distributions is that confidence intervals cannot be calculated for the parameters of the ex-Gaussian distribution functions. This is because the standard errors of the parameters are unknown. Future research could investigate whether bootstrap estimates could be used to calculate standard errors and confidence intervals. Another topic for investigation is whether it would be better to use a theoretical distribution with parameters that have known standard errors. This is the case for the Weibull distribution, which has occasionally been used for distribution analysis of response time data (e.g., Cousineau et al., 2002; Logan, 1992; Palmer et al., 2011).

In conclusion, it is widely known that response time data have an asymmetric distribution; positively skewed response time data require a distribution analysis to avoid errors of interpretation. One potential reason that response time distribution analyses are atypical is that at least 100 trials are recommended for each participant in each condition, which is far more trials than in most experiments. In the current tutorial, we demonstrated a distribution analysis technique that requires as few as 40 participants with 40 trials per condition. We hope that by decreasing the number of required trials and by providing flexible MATLAB code, we will encourage more researchers to thoroughly examine response time data.

Authors' note

Our thanks to Associate Editor, Professor Roland Pfister, Professor Josh Sandry, and Valentin Koob for their insightful critiques..

References

Blinch, J., Trovinger, C., DeWinne, C. R., de Cellio Marins, G., Ifediora, C. N., Nourollahimoghadam, M., Harry, J. R., & Palmer, T. B. (in press). Tradeoffs of estimating reaction time with absolute and relative thresholds. *Behavior Research Methods, in press*, 1–9. doi: 10.3758/s13428-023-02211-4.

Brown, S., & Heathcote, A. (2003). Qmle: Fast, robust, and efficient estimation of distribution functions based on quantiles. *Behavior Research Methods, Instruments, & Computers*, *35*(4), 485–492. doi: 10.3758/BF03195527.

Cousineau, D. (2017). Varieties of confidence intervals. *Advances in Cognitive Psychology*, *13*, 140–155. doi: 10.5709/acp-0214-z.

Cousineau, D., & Chartier, S. (2010). Outliers detection and treatment: A review. *International Journal of Psychological Research*, *3*(1), 58–67. doi: 10.21500/20112084.844.

Cousineau, D., Goodman, V. W., & Shiffrin, R. M. (2002). Extending statistics of extremes to distributions varying in position and scale and the implications for race models. *Journal of Mathematical Psychology*, *46*(4), 431–454. doi: 10.1006/jmps.2001.1399.

Cousineau, D., Goulet, M. A., & Harding, B. (2021). Summary plots with adjusted error bars: The superb framework with an implementation in R. *Advances in Methods and Practices in Psychological Science*, *4*, 1–18. doi: 10.1177/25152459211035109.

Cousineau, D., Harding, B., Walker, J. A., Durand, G., T.-Groulx, J., Lauzon, S., & Goulet, M.-A. (2023). Analyses of response time data in the same–different task. *Canadian Journal of Experimental Psychology*, *77*(2), 115–129. doi: 10.1037/cep0000301.

Cousineau, D., Thivierge, J. P., Harding, B., & Lacouture, Y. (2016). Constructing a group distribution from individual distributions. *Canadian Journal of Experimental Psychology*, *70*(3), 253–277. doi: 10.1037/cep0000069.

Harding, B., Tremblay, C., & Cousineau, D. (2014). Standard errors: A review and evaluation of standard error estimators using monte carlo simulations. *The Quantitative Methods for Psychology*, *10*(2), 107–123. doi: 10.20982/tqmp.10.2.p107.

Heathcote, A., Brown, S., & Cousineau, D. (2004). Qmpe: Estimating lognormal, wald, and weibull rt distributions with a parameter-dependent lower bound. *Behavior Research Methods, Instruments, & Computers*, *36*(2), 277–290. doi: 10.3758/BF03195574.

Heathcote, A., Brown, S., & Mewhort, D. J. K. (2002). Quantile maximum likelihood estimation of response time distributions. *Psychonomic Bulletin & Review*, *9*(2), 394–401. doi: 10.3758/BF03196299.

Heathcote, A., Popiel, S. J., & Mewhort, D. J. (1991). Analysis of response time distributions: An example using the stroop task. *Psychological Bulletin*, *109*(2), 340–347. doi: 10.1037/0033-2909.109.2.340.

Hockley, W. E. (1984). Analysis of response time distributions in the study of cognitive processes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *10*(4), 598–615. doi: 10.1037/0278-7393.10.4.598.

Hohle, R. H. (1965). Inferred components of reaction times as functions of foreperiod duration. *Journal of Experimental Psychology*, *69*(4), 382–386. doi: 10.1037/h0021740.

Lacouture, Y., & Cousineau, D. (2008). How to use MATLAB to fit the ex-gaussian and other probability functions to a distribution of response times. *Tutorials in Quantitative Methods for Psychology*, *4*(1), 35–45. doi: 10.20982/tqmp.04.1.p035.

Logan, G. D. (1992). Shapes of reaction-time distributions and shapes of learning curves: A test of the instance theory of automaticity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*(5), 883–914. doi: 10.1037/0278-7393.18.5.883.

Luce, R. D. (1986). *Response times: Their role in inferring elementary mental organization*. Oxford University Press. doi: 10.1093/acprof:oso/9780195070019.001.0001.

Miller, J. (1988). A warning about median reaction time. *Journal of Experimental Psychology: Human Perception and Performance*, *14*(3), 539–543. doi: 10.1037/0096-1523.14.3.539.

Moret-Tatay, C., Gamermann, D., Navarro-Pardo, E., & Fernández de Córdoba Castellá, P. (2018). Exgutils: A python package for statistical analysis with the ex-gaussian probability density. *Frontiers in Psychology*, *9*, 612. doi: 10.3389/fpsyg.2018.00612.

Palmer, E. M., Horowitz, T. S., Torralba, A., & Wolfe, J. M. (2011). What are the shapes of response time distributions in visual search? *Journal of Experimental Psychology: Human Perception and Performance*, *37*(1), 58–71. doi: 10.1037/a0020747.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59–108. doi: 10.1037/0033-295X.85.2.59.

Ratcliff, R. (1979). Group reaction time distributions and an analysis of distribution statistics. *Psychological Bulletin*, *86*(3), 446–461. doi: 10.1037/0033-2909.86.3.446.

Rouder, J. N., & Speckman, P. L. (2004). An evaluation of the vincentizing method of forming group-level response time distributions. *Psychonomic Bulletin & Review*, *11*(3), 419–427. doi: 10.3758/BF03196589.

Schmiedek, F., Oberauer, K., Wilhelm, O., Süß, H. M., & Wittmann, W. W. (2007). Individual differences in components of reaction time distributions and their relations to working memory and intelligence. *Journal of Experimental Psychology: General*, *136*(3), 414–429. doi: 10.1037/0096-3445.136.3.414.

Tryon, W. W. (2001). Evaluating statistical difference, equivalence, and indeterminacy using inferential confidence intervals: An integrated alternative method of conducting null hypothesis statistical tests. *Psychological Methods*, *6*(4), 371–386. doi: 10.1037/1082-989X.6.4.371.

Table A1 ■ Average Correlations Within and Between Conditions for the Three Parameters

		μ	σ	τ
Within condition	μ	1.00		
	σ	.48*	1.00	
	τ	.14	-.03	1.00
		μ	σ	τ
Between conditions	μ	.73*	.18	.34*
	σ	.15	.13	.16
	τ	.34*	.17	.42*

Note. * indicates $p \leq .05$.

Van Zandt, T. (2000). How to fit a response time distribution. *Psychonomic Bulletin & Review*, *7*(3), 424–465. doi: 10.3758/BF03214357.

Welford, A. T. (Ed.). (1980). *Reaction times*. Academic Press.

Whelan, R. (2008). Effective analysis of reaction time data. *The Psychological Record*, *58*(3), 475–482. doi: 10.1007/BF03395630.

Appendix: The Advantage of Using Multiple Participants With Few Data Points in the Geometric Quantile Averaging Framework

Suppose we have the choice between testing one participant and getting 100 observations in the condition of interest or testing ten participants and getting ten observations for each participant in the condition of interest; which would be better? Both scenarios result in a total of 100 observations. The following simulations suggest that the two scenarios are comparable when the participants have homogenous response time distributions. Further simulations will investigate the more realistic scenario of heterogeneous participants. These simulations will suggest that the geometric mean of quantiles from a moderate number of participants each having a moderate number of observations has the potential to return efficient estimates. The relationships between participants are critical to these latter simulations, and so we will begin by looking at the intercorrelations of the response time distribution functions.

Parameter Intercorrelations Within and Between Conditions

To characterize the nature of the heterogeneous response time distributions, we estimated the three parameters of the ex-Gaussian distribution function (μ, σ, and τ) for each of the forty-one participants in each of the four conditions (vision short, vision long, no vision short, and no vision long). Next, we computed the Pearson correlation coefficients between the combinations of the three parameters and the four conditions for a 12 × 12 correlation matrix. As we were not so much interested in the intercorrelations in specific conditions of this experiment, we averaged the four sets of correlations within the same conditions (a symmetrical 3 × 3 matrix) and the six sets of correlations involving different conditions (a non-symmetrical 3 × 3 matrix). Both correlation tables are shown in Table A1. With 41 participants, correlations greater than or equal to .31 are significant (two-tailed tests).

Between conditions, there was a strong correlation between μs (mean $r = .73$) and a strong correlation between τs (mean $r = .42$). As both parameters affect mean response time ($\bar{X} = \mu + \tau$), we interpret these correlations by a general speed factor; participants that are fast tend to be fast in most conditions. Similarly, τ correlates with μ of other conditions and vice-versa, reflecting that mean response time of participants manifest in both parameters. The relationships within conditions are needed to determine the structure of the heterogeneity. The result supported that the only strong correlation was between μ and σ. This suggests that slower participants tend to have more variable response time. Cousineau et al. (2023) found a similar heterogeneity of response time albeit with different methodology. These relationships will be used to model heterogeneity in our subsequent simulations.

Simulating Distributions From a Single Participant, Multiple Homogeneous Participants, and Multiple Heterogeneous Participants

The technique presented in the current tutorial is based on the quantile maximum probability estimator (QMPE; Heathcote et al., 2004). This method estimates parameters by fitting the cumulative distribution function (CDF) of the response time distributions, which is contrary to maximum-likelihood estimation (MLE) that fits the likelihood of the dataset. Van Zandt

(2000), probably the most comprehensive examination of distribution fitting, examined parameter recovery using CDF, MLE, and five other fitting techniques. She concluded that CDF and MLE fitting were the two best and noted that MLE slightly outperformed CDF fitting (see Parameter Recovery, The ex-Gaussian section on pp. 449-450). Examination of her Figure 16 indicates that CDF fitting is approximately 1.2 times more variable (less efficient) than MLE fitting when sample size is 100. As 100 is the consensual minimum sample size for fitting parameters from a single participant (Lacouture & Cousineau, 2008; Ratcliff, 1979; Van Zandt, 2000), the question is, therefore, whether 40 heterogeneous participants, each with 40 data points, will result in standard error of the CDF some 1.2 times smaller than with one participant having 100 data points. This would be comparable to the performance of MLE fitting of one participant with 100 data points.

To establish a baseline, we generated 100 simulated response time values from a single participant and estimated one point on the resulting CDF curve. The participant's response time was assumed to follow an ex-Gaussian distribution function with parameters $\mu = 500$, $\sigma = 100$, and $\tau = 250$ (Table A2, Baseline). We repeated this over 50,000 replications to get the mean and standard deviation of the estimate; the latter is the standard error. We reported the results of estimating the .50 quantile (Table A2, Simulation 1; i.e., median response time) but similar results were found for other quantiles from .05 to .95.

In a second simulation, we generated ten homogeneous participants each with ten response time values. The participants were homogeneous in that their response time followed the same distribution (ex-Gaussian with parameters $\mu = 500$, $\sigma = 100$, and $\tau = 250$). We then estimated the .50 quantile of each participant and computed the geometric mean of these estimates to emulate geometric quantile averaging (GQA). We repeated this 50,000 times and reported the mean and standard error (Table A2, Simulation 2). The estimate of the median was biased downward. This was caused by the positively skewed distribution. With small samples, the lower part of the distribution can be overrepresented, which drags the estimate of the median downward. However, the standard error was comparable to the first simulation. In other words, the efficiency of the estimates was comparable for one participant with 100 data points and from ten homogenous participants each with ten data points.

In a third simulation, we reproduced the above and introduced heterogeneity in the participants' true parameters. To that end, we varied μ across simulated participants by adding a random number between -40 and +40. The standard deviation of such random numbers is approximately 23, which matched the observed standard deviations in the estimates of μ ([22.7, 26.7]). To introduce correlation between μ and σ, we added one-fifth of the random number to σ. As before, we used the geometric mean over the estimates and reported the mean and standard error across 50,000 simulated sets of ten participants (Table A2, Simulation 3). Unsurprisingly, the standard errors increased to 34.5 because of the heterogeneity of participants.

The first three simulations were rerun with 1600 data points (i.e., a single participant with 1600 observations or 40 participants with 40 observations). Standard error was much smaller for a single participant and for 40 homogeneous participants (Table A2, Simulations 4 and 5, respectively). It was about four times smaller compared to the equivalent stimulations with 100 data point. This improvement in efficiency is predictable because we multiplied the number of data points by 16, and the standard error decreased by $\sqrt{16} = 4$. More critically, with heterogeneous participants, the standard error was 24.0 (Table A2, Simulation 6), which was slightly smaller than the baseline condition (26.6). Hence, fitting CDF from 1600 data points from heterogeneous participants is preferable to fitting CDF from 100 data points (single or homogenous participants).

As CDF was slightly less efficient than MLE (by a factor of 1.2; Van Zandt, 2000), we note that the benefit here is not 1.2 times smaller, but close (about 1.1 times smaller). Consequently, using 40 participants each with 40 response time values, we almost match the precision afforded by fitting a single participant with 100 data points. In sum, these simulations suggest that the geometric mean of quantiles from a moderate number of participants each having a moderate number of observations has the potential to return efficient estimates.

Open practices

The *Open Data* badge was earned because the data of the experiment(s) are available on osf.io/be7gp/
The *Open Material* badge was earned because supplementary material(s) are available on osf.io/be7gp/

Citation

Blinch, J., & Cousineau, D. (2024). Response time distribution analysis of medium-sized datasets in MATLAB. *The Quantitative Methods for Psychology*, *20*(2), 106–120. doi: 10.20982/tqmp.20.2.p106.

Table A2 ■ Mean and Standard Error of the Estimates for the .50 Quantile Across 50,000 Replications in Various Scenarios

Scenario	Description	Mean Estimate (ms)	Standard Error of the Mean Estimate (ms)
Baseline	True value of the .50 quantile	691	n/a
Simulation 1	One participant with 100 data points	689	26.6
Simulation 2	Ten homogeneous participants each with ten data points	670	25.2
Simulation 3	Ten heterogeneous participants each with ten data points	670	34.5
Simulation 4	One participant with 1600 data points	691	6.6
Simulation 5	Forty homogeneous participants each with 40 data points	685	6.5
Simulation 6	Forty heterogeneous participants each with 40 data points	684	24.0

Received: 01/12/2023 ~ Accepted: 23/06/2024

Listing 3 follows.

Listing 3 ■ Most of the QMPE Code From analyse.m

```
%% Quantile maximum probability estimator (QMPE) Step 1
% Join 0 and Inf as the first and last group quantiles.
quantile_geometric_mean_2_array = ones(size(quantile_geometric_mean_array,1)+2,
    condition_total) .* NaN;
quantile_geometric_mean_2_array(1,:) = 0;
quantile_geometric_mean_2_array(end,:) = Inf;
quantile_geometric_mean_2_array(1+1:end-1,:) = quantile_geometric_mean_array;

%% QMPE Step 2
% Determine the ex-Gaussian parameters (mu, sigma, and tau) with QMPE.
% QMPE requires ExGausscdf.
parameter_condition_array = ones(3,condition_total) .* NaN;
for condition_num = 1:condition_total

    % Reasonable starting search values
    tau_init = mean(std(rt_array(:,condition_num,:),0,1)) .* 0.8;
    mu_init = mean(mean(rt_array(:,condition_num,:),1)) - tau_init;
    sigma_init = sqrt(mean(var(rt_array(:,condition_num,:),1)) - tau_init^2);

    [ parameter_array, minimum_log_likelihood, return_code ] = QMPE(
    quantile_geometric_mean_2_array(:,condition_num), 'exGaussian', 'plotF', 0, '
    startPoint', pinit);
    parameter_condition_array(:,condition_num) = parameter_array';

    if (return_code ~= 1)
        error('QMPE died!');
    end
end

%% QMPE Step 3
% Calculate the ex-Gaussian probability density functions.
pdf_array = ones(1001,condition_total);
for condition_num = 1:condition_total
    pdf_array(:,condition_num) = exGaussianpdf(0:1000, parameter_condition_array(1,
    condition_num), parameter_condition_array(2,condition_num),
    parameter_condition_array(3,condition_num))';
end
```