EPA/600/R-19/104

Technical Manual: SSD Toolbox Version 1.0

Matthew Etterson

US Environmental Protection Agency
Office of Research and Development
Center for Computational Toxicology and Exposure
Great Lakes Toxicology and Ecology Division
Duluth, MN

March 2020

Contents

Glossary

Term	Definition
AIC	Akaike Information Criterion
AIC_c	Akaike Information Criterion adjusted for small sample size
EC10	Concentration expected to cause an effect in 10% of test subjects
ECDF	Empirical Cumulative Distribution Function
GUI	Graphical User Interface
GoF	Goodness-of-fit
HC05	Concentration expected to be hazardous to 5% of species tested
HCp	Concentration expected to be hazardous to p% of species tested
Hessian Matrix	Matrix of second derivatives of the Log-likelihood at the MLE
LC50	Concentration expected to be lethal to 50% of test subjects
LD50	Dose expected to be lethal to 50% of test subjects
MCMC	Markov Chain Monte Carlo
MCR	Matlab Compiler Runtime
mg a.i./kg bw	Milligrams active ingredient per kilogram body weight
MLE	Maximum Likelihood Estimator
SSD	Species Sensitivity Distribution

Introduction

Species sensitivity distributions are a common tool used for setting safe limits on chemical concentrations in surface waters (Posthuma et al. 2002, Suter 2002, Chapman et al. 2007, TenBrook et al. 2010). Although the analysis and interpretation of species sensitivity distributions varies widely, the basic methodology is quite general and can be summarized as a three-step procedure. First, results from separate toxicity tests on a given chemical using several species are compiled. Second, a statistical distribution to which the test results are thought to conform is chosen and fit to the data. Third, the fitted distribution is used to infer a concentration that will be protective of a desired proportion of similar species for which inference is desired.

The procedure described above necessarily relies upon policy decisions, for example, concerning the proportion of species that should be protected and the necessary level of

confidence with which the protective concentration is identified. This manual focuses on procedures for fitting statistical distributions and providing risk assessors with the information to assess the quality of a fitted distribution. As such its content concerns the three steps laid out in the previous paragraph, each of which can be accomplished in different ways, with different results. This manual does not consider the underlying policy decisions required for application of species sensitivity distributions to regulatory decision making. This manual is intended to be a companion to the SSD Toolbox User's Guide, which gives step-by-step instructions on how to use the software.

Data

An important consideration in data compilation is the distribution of available data across species. Strictly speaking, an assumption of all the methods considered below is that the data (test results) pertain to a random sample of species from the species group for which the analysis is intended to apply. This assumption is always violated; a relatively limited subset of taxa makes up the greater part of all toxicity tests. Therefore, in using SSDs to derive protection goals, one should consider the potential biases in the data set relative to the group for which the protection goal is intended to apply. Another important, and related, consideration concerns the independence of data. Again, strictly speaking, an assumption underlying most SSD methods is that the data are independent and identically distributed, but this assumption may be violated if closely related species (with a similar toxicity response) are included in the data (Moore et al. 2019).

In most cases, the set of data used for fitting an SSD will require considerable curation prior to distribution fitting. While the specific steps required will differ among applications and datasets, some important common steps should be employed:

1) Endpoints should be commensurate across species (e.g., all endpoints are LC50s, or all endpoints are EC10s).

2) All endpoints should be expressed in the same units (e.g., mg a.i./kg bw).

3) Endpoints should be derived from studies with similar designs (e.g., similar exposure durations).

Other considerations may also apply, at the user's discretion. The goal of such steps is to ensure that the variation observed represents true variation among species, where confounding factors are controlled to the greatest extent possible. However, exclusion of test results due to data curation will reduce sample size, resulting in greater uncertainty (sampling variance & confidence limits) in the desired inferential endpoint (*e.g.*, the HC05), but potentially less bias.

Transformations and Body-weight scaling

When data are imported to the SSD Toolbox, several transformations may be performed in the following order. First, when the body-weight scaling function is used, the individual toxicity endpoints are first rescaled according to the body weight scaling function provided by Mineau et al. (1996):

Eq. (1) $$Scaled\ LD50 = LD50\left(\frac{Target\ Weight}{Tested\ Weight}\right)^{(x-1)}$$

In Eq. (1), *LD*50 is the unscaled LD50 resulting from the toxicity test, *Tested Weight* is the mean weight of tested birds, and *Target Weight* is the weight to which the toxicity data are intended to be standardized (default = 100g). The resulting Scaled LD50s are used for all subsequent analyses. Second, when there are multiple toxicity values for a given taxon the geometric mean of those values is calculated. This applies equally to unscaled and scaled (i.e., as per step one, above) toxicity values. Finally, when the normal, logistic, triangular, or Gumbel distributions are fit to the geometric mean toxicity values, those values are first common log (log_{10}) transformed. The resulting HC05 estimates are provided on the natural scale.

Choosing and fitting a distribution

Many statistical distributions have been used for fitting SSDs (*e.g.*, log-normal, log-logistic, $Burr_{III}$, etc.); however, several analyses have shown that no one distribution is preferred across datasets (Newman et al. 2000, Zajdlik & Associates 2005, Chapman et al. 2007). Deciding which statistical distribution to fit to a set of data has been described as one of the most important and difficult choices in the use of species sensitivity distributions (Chapman et al. 2007). Two important additional decisions must be made when fitting a distribution to empirical data. The first concerns how the distribution will be fit to the data, which is equivalent to the problem of parameter estimation. The second choice concerns how to assess the quality or accuracy of the fitted distribution as a general representation of the data, or goodness-of-fit. A related concern involves deciding how to choose among the fitted distributions when multiple distributions are fit to the same data. For many types of distributions (*e.g.,* normal, logistic, triangular, Gumbel) data are usually transformed prior to analysis, most frequently using the common log (log_{10}) transformation. This complicates comparisons between distributions fit to transformed versus untransformed data.

Four methods commonly used for estimating the parameters of SSDs are implemented in the SSD Toolbox. These are maximum likelihood, moment estimators, graphical methods, and Bayesian methods (specifically the Metropolis-Hastings algorithm). These methods are described in more detail below. Not all methods can be used with all distributions.

Newman et al. (2000) recommended a non-parametric method for fitting SSDs using empirical bootstrapping. However, given the common regulatory interest in the fifth percentile of acute values, bootstrap estimation does not seem feasible because it would require at least 19 data points to estimate the fifth percentile of the empirical cumulative distribution function (ECDF). Bootstrap methods are used below to test goodness-of-fit and to estimate sampling variance.

Chapman et al. (2007) emphasized the importance of visual inspection of fitted distributions against the empirical data to which they are fit. With small sample sizes, visual inspection may be the most reliable method for assessing fit, despite its obvious subjectivity. The SSD toolbox emphasizes data and curve visualization to facilitate such inspection. More details on estimating posterior goodness-of-fit are provided below.

Maximum likelihood

Maximum likelihood methods for SSDs were first tested by Kooijman (1987) who reported substantial bias in estimation of the scale parameter (β) for the logistic distribution with sample sizes ≤ 5 (logistic formulae given below). Shao (2000) described maximum likelihood estimators for the $Burr_{III}$ distribution; these estimators are implemented by the software BurrliOZ (Campbell et al. 2000). Several recent reports on the application and analysis of species sensitivity distributions have employed maximum likelihood with other distributions (Zajdlik and Associates 2005, Chapman et al. 2007), often citing the first of the following desirable properties. First, when data fit the assumed distribution, maximum likelihood parameter estimators (MLEs) are the most efficient parameter estimators possible (*i.e.*, the estimators that produce the smallest sampling variance, Edwards 1992), though they may be biased. Second, the use of maximum likelihood allows the fit of different distributions to be compared using information theoretic methods for comparing models (Burnham and Anderson 2002). Third, use of maximum likelihood allows model-averaging of estimated quantiles, such as the HC05 (Burnham and Anderson 2002) across multiple distributions. Fourth, maximum likelihood and restricted maximum likelihood (Harville 1977), allow specification of hierarchical models that may otherwise be very difficult to fit.

The SSD Toolbox formulates the log-likelihood equations for each of the five distributions as the natural logarithm of the probability density function (f) for that distribution. The resulting log-likelihood is summed over all data points:

Eq. (2) $L(\boldsymbol{\theta}|\boldsymbol{X}) \propto \sum_{i=1}^{n} ln\big(f(x_i|\boldsymbol{\theta})\big)$

Maximum likelihood estimates (MLEs) of distribution parameters (here represented as a vector, $\boldsymbol{\theta}$) are those values that, when substituted into Eq. 2 maximize its value. These are found by numerical search, which can be slow. Note that in some cases (e.g., normal distribution) closed-form arithmetic expressions are available for the MLEs. However, the SSD Toolbox does not make use of closed-form estimators, in part to force all maximum likelihood estimates to be

obtained in the same way, and in part to take advantage of the numerical estimation of the Hessian matrix, which is useful for estimating the sampling variance of the parameters and the hazardous concentrations (HCp).

Moment Estimators

Moment estimators are a common method for fitting SSDs (Kooijman 1987, Van Straalen and Deneman 1989). In practice, they work by equating the mean and variance of a sample to the parametric mean and variance of a chosen distribution, which are functions of the parameters of that distribution. This creates two equations in two unknowns, which can then be solved for the unknown parameters. The resulting solution is an estimate of the parameters of the distribution expressed as functions of the sample mean and sample variance. Although this procedure has been described in terms of the mean and variance (the first two moments), it could be extended to higher moments as well if a distribution (*e.g.*, Burr_{III}) has more than two parameters.

Moment estimators were derived, wherever possible, by setting the expected distributional mean and variance equal to the sample mean and variance and solving for the distributional parameters. For example, let $\bar{x}$ and s^2 represent the sample mean and variance, respectively (regardless of assumed distribution). The mean and variance of a logistic distribution are α and $\frac{\pi^2}{3}\beta^2$. Setting $\alpha = \bar{x}$ and $s^2 = \frac{\pi^2}{3}\beta^2$ and solving for α and β results in the two moment estimators $\hat{\alpha} = \bar{x}$ and $\hat{\beta} = \frac{s}{\pi}\sqrt{3}$, where the circumflex over the parameter symbols indicates that they are estimated quantities. For the Burr_{III} distribution, moment estimators were not derived because the Burr distribution has three parameters, requiring three equations, but an equation for the third moment of the Burr_{III} distribution was not immediately available. Moment estimators for four distributions (normal, logistic, triangular, Gumbel) are included in the SSD Toolbox and presented below.

Linearization

Linearization, or graphical methods, for use in SSDs was described by Erickson and Stephan (1988). Linearization is a subset of the general theory of order statistics (Arnold et al. 2008) and has two unique attributes that make it attractive for use in fitting SSDs (TenBrook et al. 2008). First, once data have been ordered and the empirical percentiles obtained, the linear estimation model can be weighted toward the lower tail of the distribution, which is generally the portion of the distribution of interest for regulation and risk assessment. This can alleviate potential biases resulting from skewed toxicity distributions. Second, and related to the first, toxicity test results that are right-censored (known only to be greater than the highest tested concentration or dose) can often be accommodated (Erickson and Stephan 1988). Linearization can be used on any distribution for which the cumulative distribution function can be linearized

through transformation. Of the six distributions in the SSD Toolbox, the normal, logistic, triangular and Gumbel and Weibull can be fit using graphical methods.

Graphical estimation is implemented in the SSD Toolbox using a linearization of the cumulative distribution function or a standard form of the distribution (parameters chosen so that mean = 0, variance = 1). For example, a normal distribution can be standardized (i.e., to z scores) as $z = \frac{y-\mu}{\sigma}$, where y = $\log_{10}$ toxicity value, and μ and σ are the usual parameters of the normal distribution. The z-scores are the quantiles of the standard normal distribution. Rearranging this equation gives:

Eq. (3) $y = \sigma z + \mu$

Equation (3) is a linear function with slope σ and intercept μ. Given paired values of z and y, σ and μ can be estimated by linear regression. Importantly, σ and μ can be estimated from any subset of ordered pairs of z_i and y_i, such as the lower 50% of values. It should be noted, however, that the standard error of σ and μ will grow as the quantile is lowered, because the linear regression (Eq. 2) will include fewer data points.

Paired values of z and y for use in Eq. (2) are obtained using the empirical cumulative distribution function (ECDF) of y, which gives the cumulative probability associated with each value in y. Using these probabilities, the standard scores for the desired distribution can be obtained from the inverse cumulative distribution function (F^{-1}) for the standard form of the chosen distribution. For the normal distribution, these are typically referred to as z-scores. A similar linearization procedure, with some variation in details, is followed for other distributions fit using graphical methods.

The empirical cumulative probabilities (p) for the ECDF in the SSD Toolbox are calculated for the i^{th} variate in y as:

Eq. (4) $p_i = \frac{r_i}{n+1}$

In Equation (4) r_i is the rank of the i^{th} variate in y and n is the number of variates (species for which toxicity test results are available). Alternative choices for calculating the p_i (sometimes referred to as plotting points) exist in the literature (reviewed by Erickson and Stephan 1988), however, the SSD Toolbox has adopted Eq. (3) in part because it corresponds to quantiles for the well-defined ECDF. Alternative plotting positions would result in different estimates of the hazardous concentrations.

Bayesian Methods

Bayesian methods, like maximum likelihood, rely on the likelihood function for the distribution parameters given the data. However, Bayesian estimation also incorporates existing knowledge in the form of prior distributions on model parameters (i.e., the distribution parameters in the

SSD context). Bayesian methods work by sampling from the posterior distribution of the parameters, conditional on the priors and the likelihood evaluated on the data (King et al. 2010, Link et al. 2010). In rare cases the posterior distribution is analytically tractable, but those cases are not considered here. Often posterior distributions are sampled using a Markov Chain Monte Carlo algorithm (MCMC, Link et al. 2010). The specific MCMC algorithm implemented in the SSD Toolbox is the Metropolis-Hastings algorithm (Hastings 1970). In the current version, the SSD Toolbox employs vague priors (uniform over the range of potential parameter values). A future version will allow greater user control over prior distributions.

Goodness-of-fit

Goodness-of-fit is a measure of how well an assumed distribution fits a set of data, given the data and the values of the estimated parameters of the distribution. Numerical tests for goodness-of-fit can be divided into parametric and non-parametric methods. In either case, the test begins with the definition of a test statistic that can be reliably predicted to increase in magnitude with lack of fit. In the SSD Toolbox the discrepancy statistic is the sum of the squared differences between the percentiles of the ECDF and the cumulative distribution function (F) for the fitted distribution. With parametric goodness-of-fit tests, a theoretical distribution for the test statistic can be derived, and probabilities are estimated from that theoretical distribution. The derivation of the theoretical distribution of the test statistic often depends on the hypothesized distribution for the data. Therefore, parametric goodness-of-fit tests tend not to work well at small sample sizes and generally apply to only one distribution (often the normal distribution). In contrast, non-parametric methods often work by statistical resampling methods (Efron and Tibshirani 1994) and probabilities are assessed as simple ranks of observed statistics among a set of simulated statistics. They can be applied to any continuous distribution and are valid regardless of sample size. However, both parametric and non-parametric methods lack power at small sample sizes. Thus, for sample sizes typically available for SSD analysis, visual inspection may be a more reliable method for diagnosing lack-of-fit than numerical analyses.

Luttik and Aldenberg (1997), Aldenberg and Luttik (2002), and Newman et al. (2000) all considered parametric goodness-of-fit tests. Zajdlik & Associates (2005) recommended the Anderson-Darling test for all distributions, except the normal and log-normal distributions for which they recommended the Shapiro-Wilks test. Chapman et al. (2007) carried out extensive simulations of the power properties of goodness-of-fit tests for the normal distribution and concluded that power to detect non-normality (lack of fit) was extremely low, especially at sample sizes < 20.

Newman et al. (2000) and Shao (2000) employed non-parametric goodness-of-fit tests based on empirical bootstrap sampling (Efron and Tibshirani 1994, Manly 1997). Chapman et al. (2007)

described a parametric bootstrap procedure (also described by Efron and Tibshirani 1994) but did not apply it to goodness-of-fit testing for SSDs. Because of their utility at all sample sizes, and applicability to all continuous distributions, only bootstrap methods are implemented in the current version of the SSD Toolbox, however parametric tests may be added to a future release.

Numerical methods

The SSD Toolbox uses bootstrap sampling to generate replicate sets of data based on the data under analysis. The process begins after a distribution is fit to the data and the discrepancy statistic described above is calculated (the sum of squared distances between the empirical and parametric cumulative distribution functions). New data sets, of the same size as the original data, are generated by drawing random samples from the fitted distribution (parametric bootstrap). These random samples represent plausible data sets, of the same size as the original data, that could be observed if the distribution truly fits the original data. To these new data sets the same distribution is fit and the discrepancy statistic is calculated for the simulated data under the newly fitted distribution. This process is repeated a specified number of times (the SSD Toolbox default is 1,000 bootstrap samples) to generate a distribution of test statistics that would be expected if the data were drawn from the fitted distribution. Large values of the discrepancy statistic indicate poorer fit, whereas small values indicate better fit. The proportion of simulated discrepancy statistics that are greater than or equal to the observed discrepancy statistic for the empirical data is interpreted as the P-value for lack of fit. Small P-values indicate that the discrepancy statistic for the empirical data is larger than most of the simulated values, suggesting that the distribution fits the empirical data more poorly than would be expected by chance.

The procedure described immediately above also generates a distribution of parameter values that can be used to estimate sampling variance, which is described in greater detail below.

Visual inspection

At sample sizes typically available for fitting SSDs in ecotoxicology (often less than 20), numerical methods for assessing fit will suffer from low statistical power, resulting in poor ability to identify lack of fit (Chapman 2007). Therefore, visual inspection is an important step in acceptance/rejection of a candidate distribution for a dataset and for deciding what kind of inference should be made from the fitted distribution. However, visual inspection is necessarily a subjective exercise and should be used cautiously and transparently. Below I provide some examples using permethrin LC50 data for a variety of aquatic taxa taken from Fojut et al. (2012) and chlorpyrifos LC50 data for aquatic invertebrates taken from USEPA 2016.

Upon import, the SSD Toolbox gives you the choice to view histograms of sampling density (number of toxicity values per taxon), toxicity (number of geometric mean values in binned

ranges of toxicity) and a plot of the ECDF. The latter two are intended for use in thinking about fitting a distribution. For example, Figures 1 and 2, below show sample datasets that are both skewed, but in opposite directions, suggesting that different distributions may be required to fit the two respective datasets.

Figure 1. Histogram of $\log_{10}$ LC50s for invertebrates exposed to Chlorpyrifos

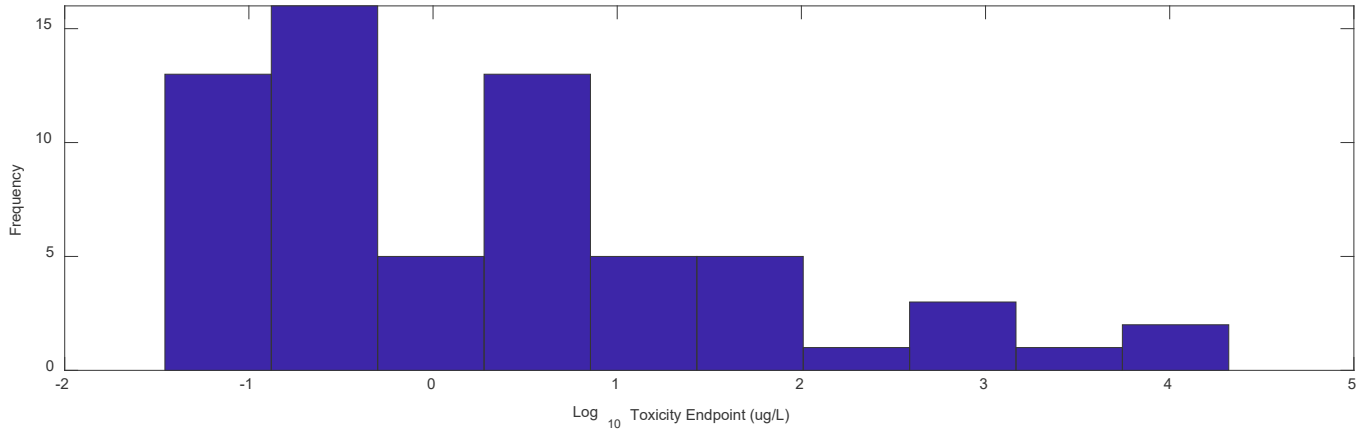

In Figures 2 and 3, below, the chlorpyrifos invertebrate data are fit using a Gumbel distribution (Fig. 2), and a Weibull distribution (Fig. 3). Although model selection criteria (e.g., AIC_c, BIC) would objectively score the Gumbel higher, the superior performance of Gumbel is easily seen by visual inspection.

Figure 2. Gumbel Cumulative Distribution Function for $\log_{10}$ LC50s for aquatic invertebrates exposed to Chlorpyrifos

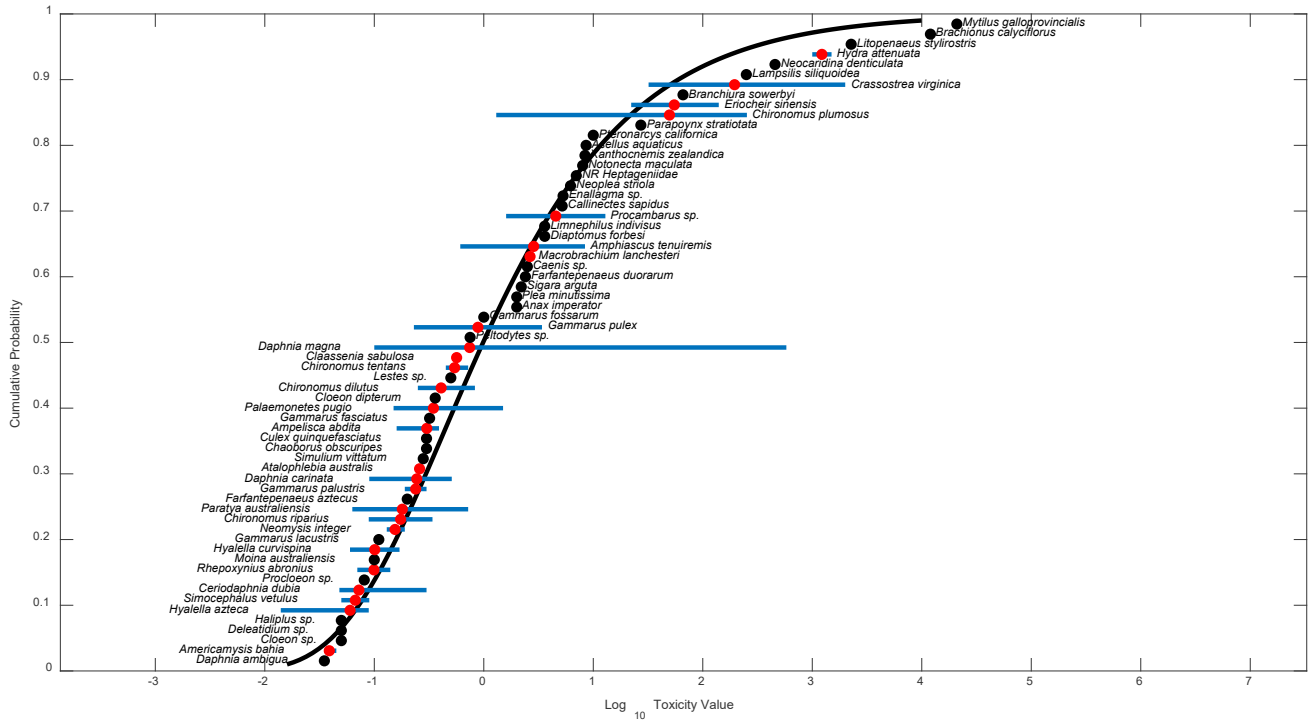
Cumulative Probability
Log 10 Toxicity Value
Mytilus galloprovincialis
Brachionus calyciflorus
Litopenaeus stylirostris
Hydra attenuata
Neocaridina denticulata
Lampsilis siliquoidea
Crassostrea virginica
Branchiura sowerbyi
Eriocheir sinensis
Chironomus plumosus
Parapoynx stratiotata
Pteronarcys californica
Asellus aquaticus
Xanthocnemis zealandica
Notonecta maculata
NR Heptageniidae
Neoplea striola
Tenaliagma sp.
Callinectes sapidus
Procambarus sp.
Limnephilus indivisus
Diaptomus forbesi
Amphiascus tenuiremis
Macrobrachium lanchesteri
Caenis sp.
Farfantepenaeus duorarum
Sigara arguta
Plea minutissima
Anax imperator
Gammarus fossarum
Gammarus pulex
Tubifex tubifex
Daphnia magna
Claassenia sabulosa
Chironomus tentans
Lestes sp.
Chironomus dilutus
Cloeon dipterum
Palaemonetes pugio
Gammarus fasciatus
Ampelisca abdita
Culex quinquefasciatus
Chaoborus obscuripes
Simulium vittatum
Atalophlebia australis
Daphnia carinata
Gammarus palustris
Farfantepenaeus aztecus
Paratya australiensis
Chironomus riparius
Neomysis integer
Gammarus lacustris
Hyalella curvispina
Moina australiensis
Rhepoxynius abronius
Procloeon sp.
Ceriodaphnia dubia
Simocephalus vetulus
Hyalella azteca
Haliplus sp.
Deleatidium sp.
Cloeon sp.
Americamysis bahia
Daphnia ambigua

Figure 3. Weibull Cumulative Distribution Function for LC50s for aquatic invertebrates exposed to Chlorpyrifos

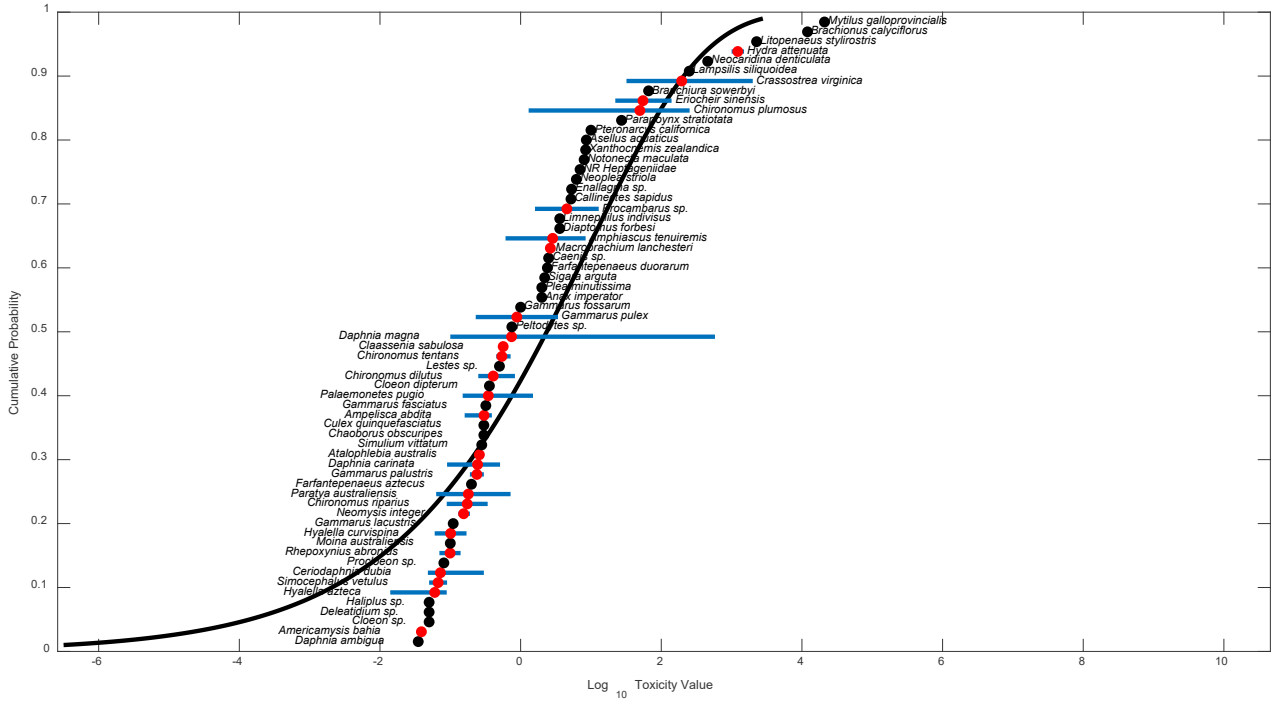

Q-Q plots can also be very useful for diagnosing lack of fit. In Figures 3 and 4, below, Q-Q plots for Gumbel and Weibull are is shown for the same Chlorpyrifos invertebrate data. In the Q-Q plots, the horizontal axis gives the empirical quantiles and the vertical axis gives the predicted quantiles (from the fitted distribution). If the model fit the data perfectly, the empirical and predicted quantiles would be the same (solid line). Figure 4 shows relatively good fit. In contrast, the Q-Q plot for the Weibull distribution fit to the same data (Fig. 5) reveals poor fit, especially in the lower tail of the distribution.

Figure 4. Normal Q-Q plot for $\log_{10}$ LC50s for invertebrates exposed to Chlorpyrifos.

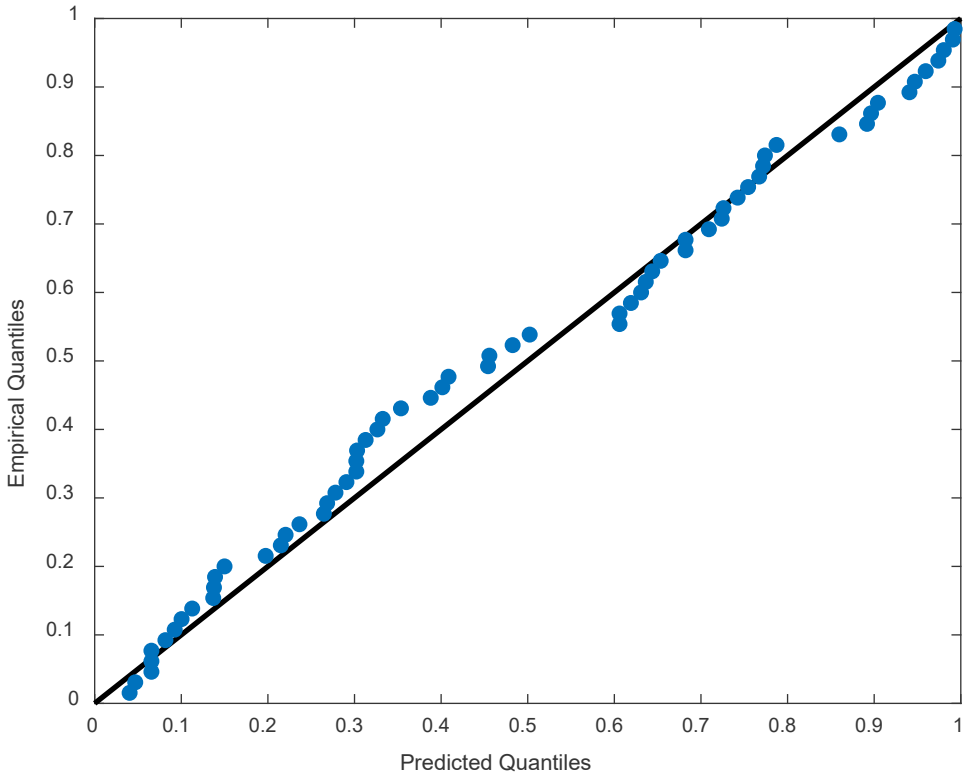

Figure 5. Normal Q-Q plot for $\log_{10}$ LC50s for invertebrates exposed to Chlorpyrifos

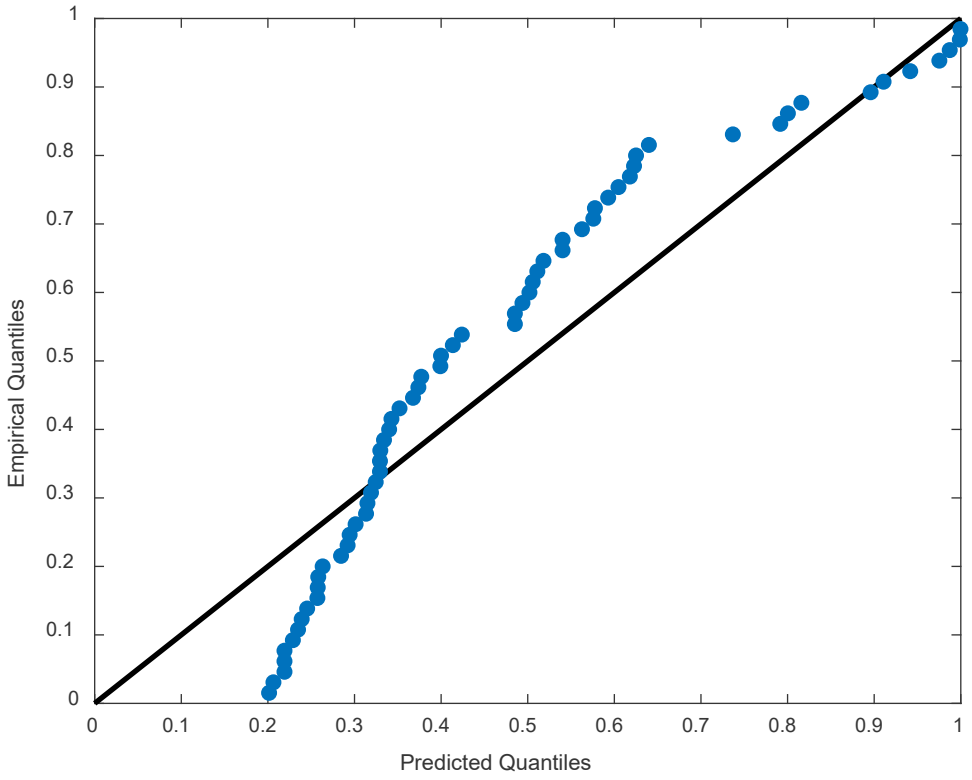

Assessing Bayesian fits using posterior diagnostics

The functions provided for evaluating distributions fit using the Metropolis-Hastings algorithm are standard methods from the Bayesian literature (King et al. 2010). Like visual goodness-of-fit evaluations described above for other methods, many Bayesian goodness of fit methods involve inspection of graphical output, looking for evidence for poor fit. As above, this necessarily involves some degree of subjectivity. Below, these methods are reviewed to highlight patterns indicating lack of fit.

Bayesian *p*-values are a *posterior* measure of fit. As calculated here, they compare a discrepancy statistic calculated on hypothetical data generated from parameter sets from the posterior distribution to the same statistic calculated for the empirical data using expected (marginal) parameter values calculated over the full posterior distribution (Link & Barker 2010). In general, the discrepancy statistic for the observed data and expected parameter values should not be markedly larger (or smaller) than the set of values calculated using specific iterations of the Markov chain Monte Carlo (MCMC). The SSD Toolbox calculates this statistic based on 10,000 random samples from the posterior distribution. The discrepancy statistic used

by the SSD Toolbox is the sum of the squared distances between the empirical and fitted cumulative distribution functions (i.e., the same function used by the bootstrap goodness-of-fit routine described above). In general, the larger the deviation of the Bayesian *p*-value from 0.5, the greater the indication of lack of fit. In other words, both large and small Bayesian *p*-values indicate lack of fit. However, I am aware of no "rule of thumb" indicating a threshold Bayesian *p*-value for rejection of a distribution. In Figure 6, below, the Bayesian P-values suggest that only the Gumbel and $Burr_{III}$ distributions are competitive for the Chlorpyrifos invertebrate data, in good agreement with the maximum likelihood analysis above.

Figure 6. SSD Toolbox output table with Bayesian *p*-values for distributions fit to the Chlorpyrifos invertebrate data

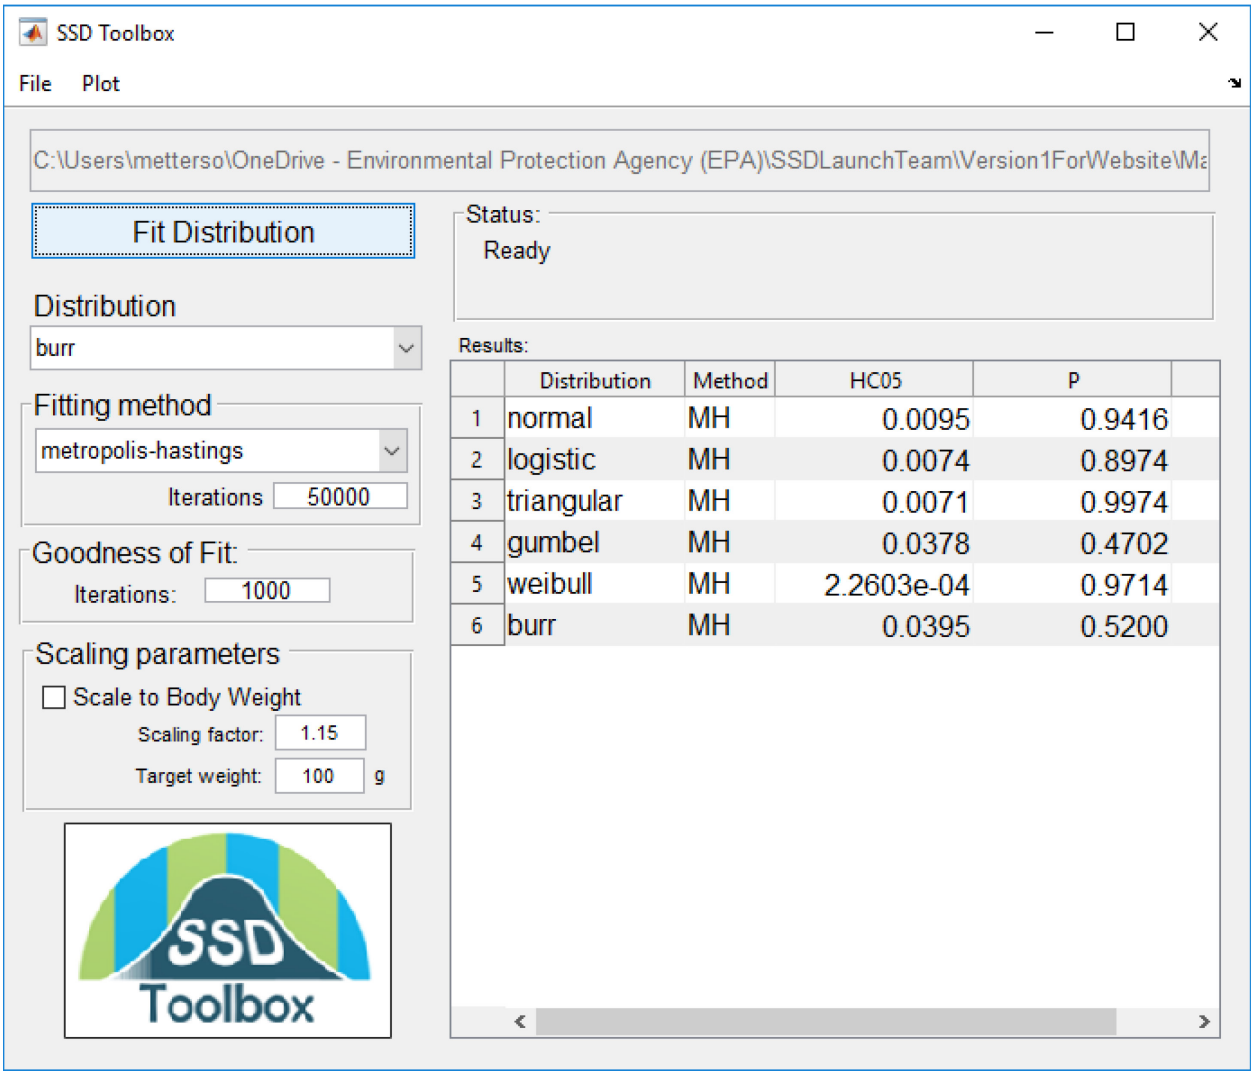

	Distribution	Method	HC05	P
1	normal	MH	0.0095	0.9416
2	logistic	MH	0.0074	0.8974
3	triangular	MH	0.0071	0.9974
4	gumbel	MH	0.0378	0.4702
5	weibull	MH	2.2603e-04	0.9714
6	burr	MH	0.0395	0.5200

Autocorrelation plots describe the relative independence among sequentially sampled points in the posterior distribution. A high degree of autocorrelation may indicate that the posterior

distribution is poorly sampled. Below, autocorrelation plots for the Gumbel distribution fit to the Chlorpyrifos invertebrate data are shown (Fig. 7).

Figure 7. Sample autocorrelation plots for the posterior distributions of Gumbel parameters (μ, β), showing decay of autocorrelation among sequential points

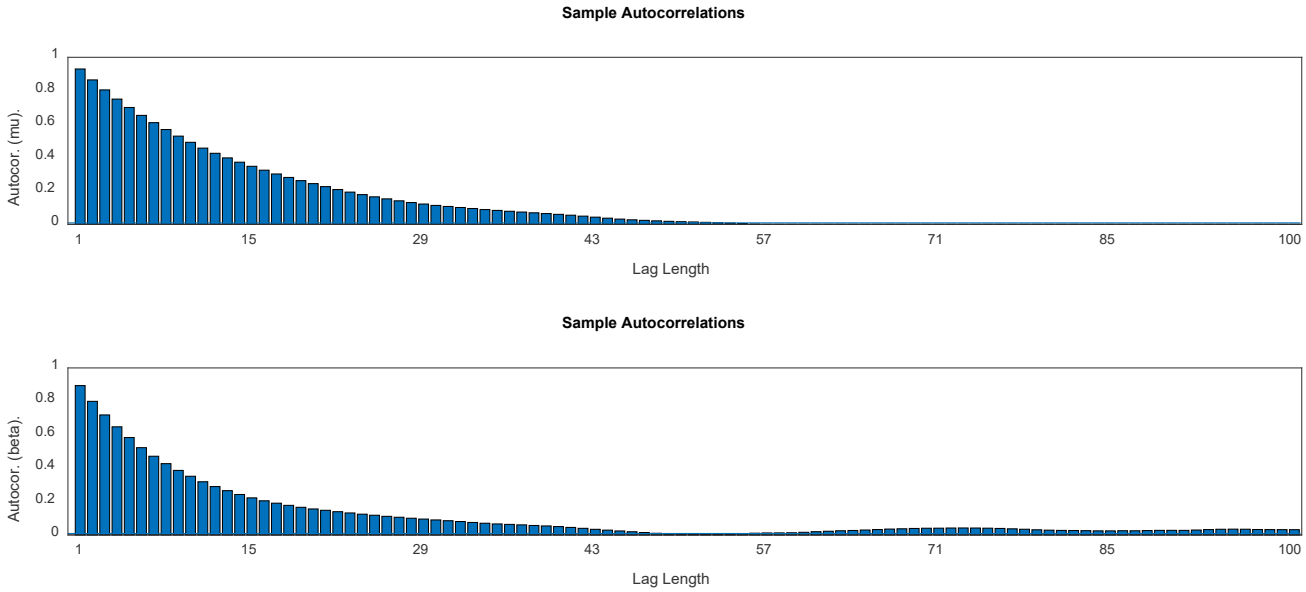

Trace plots graphically display the actual sequence of values sampled from the posterior distribution. Figure 8 shows the sequences sampled for the Gumbel distribution fit to the Chlorpyrifos data.

Figure 8. Trace plots for the Gumbel distribution fit to the Chlorpyrifos data

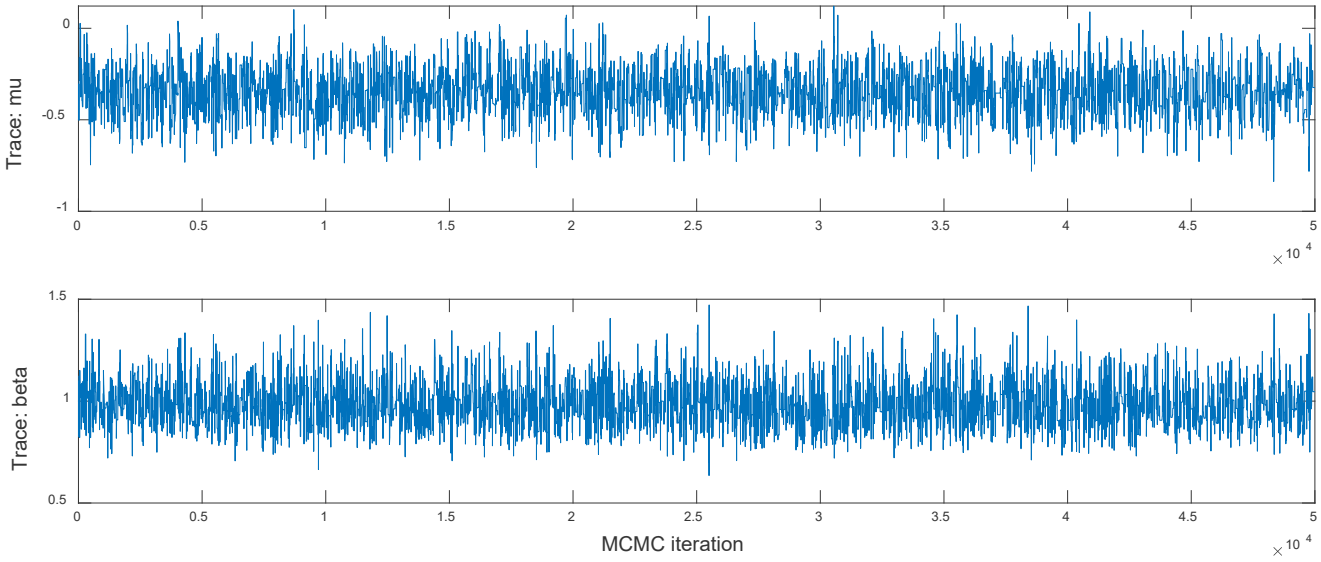

Posterior distribution plots offer a final graphical diagnostic for Bayesian model fits. Figure 9 shows the posterior marginal distributions (diagonal) for each parameter and the joint posterior distributions for each parameter pair (off-diagonal) for the Gumbel distribution fit to the Chlorpyrifos data. The densities are fairly smooth, and the joint distributions suggest little or no covariance to be concerned about (unstructured scatter-plots of joint values).

Figure 9. Posterior parameter distributions for Gumbel distribution fit to the Chlorpyrifos data

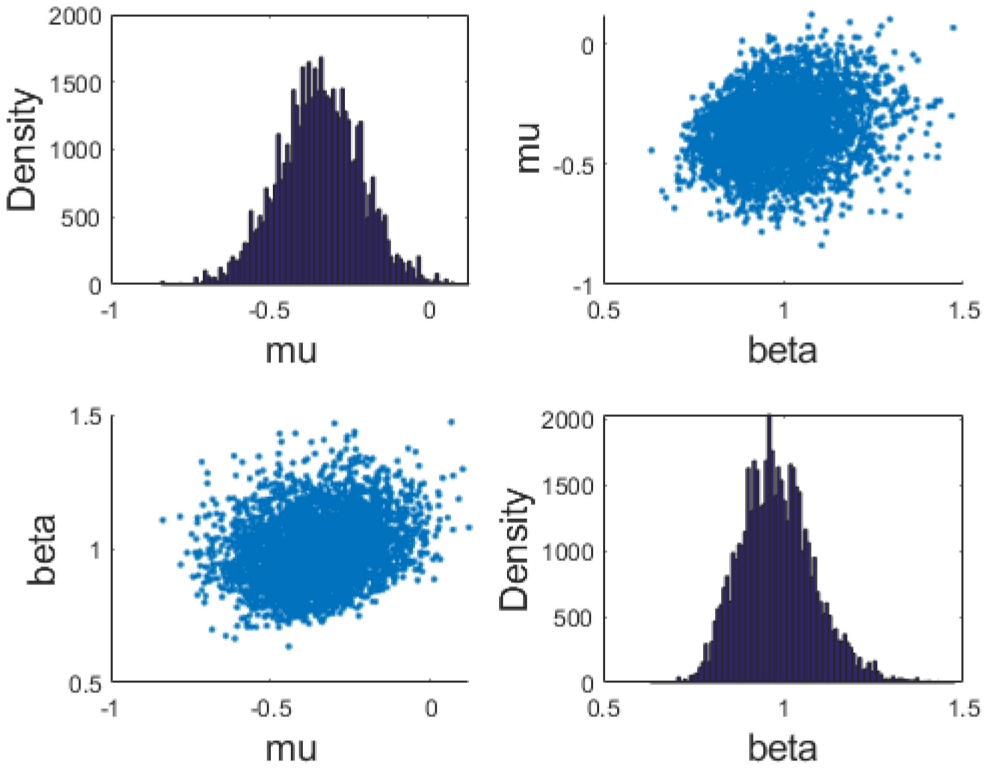

Uncertainty in the fitted distribution

All analyses of parametric species sensitivity distributions begin by estimating the parameters of the distribution (see above). Thus, the distributional parameters are a universal inferential endpoint (excluding non-parametric SSDs, Newman et al. 2000). Once the parameters are estimated, a given percentile (p) of the distribution is often chosen to represent the concentration at which (no more than) p% of species will be at risk of adverse effects, referred to herein as the HCp. Regardless of how the distribution is fit, an HCp is easily estimated using the quantile function for the fitted distribution. However, estimates of percentiles are subject to bias (if the distribution doesn't fit the data very well) and uncertainty (especially when the number of test results are limited). Methods for handling these aspects of distribution fitting vary widely in the SSD literature. Erickson and Stephan (1988) also pointed out that, by Jensen's inequality, an unbiased estimator for the HCp, might be a biased estimator of the percentile (intended to be p) of species protected at the estimated HCp if the quantile function is non-linear in p (as is generally the case).

Sampling variance and standard error

Distributional parameters estimated from empirical data are subject to sampling variance. In other words, given data that conform to a specified distribution, if equal sized (but different) sets of data are drawn from the same distribution, the parameter estimates will differ with each set of data, resulting in a distribution of parameter estimates. The variance of this theoretical distribution of parameter estimates is termed the sampling variance of the parameter estimates. If the estimation procedure is unbiased the average of the parameter estimates will be arbitrarily close to the 'true' values as the procedure is repeated more and more times. However, the expected variance in these parameter estimates may be quite large and is generally inversely related to the size of the data sample. This sampling variance is present in all four fitting techniques described above (maximum likelihood, moment estimators, linearization, and Bayesian methods), though it may differ among techniques. Sampling variance of parameter estimates translates directly (though not necessarily linearly) into sampling variance of quantiles of a distribution (*i.e.*, the estimated HC05). Common methods for estimating sampling variance around quantiles in an SSD include the delta method (Seber 1982, Shao 2000) and the bootstrap (Newman et al. 2000). The standard error of the HCp is the square root of the sampling variance of the HCp.

The SSD Toolbox offers three methods for calculating the sampling variance (and therefore standard error) of the quantiles of fitted distribution. If the distribution has been fit using maximum likelihood then the covariance matrix of the distribution parameters (e.g., μ and σ from a normal distribution, etc.) are calculated from the negative inverse Hessian Matrix (matrix of second derivatives of the log-likelihood evaluated at the MLE). From the covariance matrix, the sampling variance and standard error of the quantiles are estimated with the delta method. However, the negative inverse Hessian only asymptotically converges on the true covariance matrix and may be unreliable at small sample sizes. Therefore, the sampling variances may also be estimated using parametric bootstrap sampling (described in Goodness-of-fit section). This method also has the advantage of being available for fitting methods other than maximum likelihood. Finally, when a distribution is fit using Metropolis Hastings, the sampling variances of percentiles are calculated from the posterior distribution of parameters.

Confidence limits

Confidence limits for an estimated hazardous concentration (HCp) are an alternative expression of uncertainty in the model parameter estimates. These can be one-sided expressions of confidence that the true HCp is greater than a specified concentration (Kooijman 1987, van Straalen and Deneman 1989, Aldenberg and Slob 1993) or two-sided limits related to the probability that the region defined by a lower and upper bound would contain the true HCp (Shao 2000, Newman et al. 2000). In the SSD Toolbox, confidence limits are calculated using three different methods, depending on the method used to fit the distribution.

When maximum likelihood is used, the covariance matrix for the distribution parameters (e.g., μ, σ from a normal distribution) is available as a byproduct of the estimation routine (i.e., as the negative inverse of the Hessian matrix, as described above). Using the estimated covariance matrix, the delta method is used to calculate the sampling variance of a percentile (e.g., the HC05). The sampling variance of the percentile is then used to calculate the confidence limit around the hazardous concentration using the z-score corresponding to a 95% confidence level (z = 1.96).

For moment estimators and linearization, the sampling variance must be estimated using parametric bootstrapping. This is done using the goodness-of-fit algorithm (see above), from which the sampling distribution of the hazardous concentration (HC) is estimated. The samples from the sampling distribution of the HC are ordered and then used to calculate percentiles from the sampling distribution. With the ordered sample and corresponding percentiles, the central 95% of the distribution is estimated by finding the values corresponding to the lower 2.5% and upper 97.5% of the sampling distribution. When values do not correspond exactly to the desired percentile, the software conservatively choses the outer values. For example, the lower CL is calculated as the largest value from the sampling distribution with a corresponding percentile that is less than or equal to the desired percentile (e.g., 2.5), and similarly (greater than or equal to) for the upper CL. Parametric bootstrapping may also be used with distributions fit using maximum likelihood, if desired.

When distributions are fit using the Metropolis Hastings algorithm, 95% Bayesian Credible intervals are calculated from the posterior distribution for each quantile.

Model selection & multidistributional inference

Many researchers have discussed the important (and difficult) choice of which distribution to employ for an SSD (Newman et al. 2000, Zajdlik & Associates 2005, Chapman et al. 2007). Assessing the goodness-of-fit (see section above) of a distribution provides only limited information for comparing distributions because discrepancies of fit will generally decrease monotonically with increasing number of estimated parameters. Yet an over-parameterized model may have poor predictive ability. Formal model selection criteria impose a penalty for each estimated parameter, which creates a tradeoff between parsimony and fit. The fact that most species sensitivity distributions have two estimated parameters (though the $Burr_{III}$ distribution has three) alleviates this concern somewhat. However, model selection methods are also useful for ranking the performance of alternative distributions and for formally averaging model predictions when multiple models are fit (Burnham and Anderson 2002).

AIC, AIC$_c$, & BIC

Elphick (2011) used Akaike's Information Criterion (AIC) to compare several candidate distributions, including log-normal, log-logistic, log-Gumbel and Weibull and report similar performance. AIC may be used only when distributions are fit using maximum likelihood. The equation for AIC is:

Eq. (5) $\mathrm{AIC} = -2L + 2K$

In Eq. (5), L is the maximized log-likelihood function, and K is the number of parameters estimated in fitting the distribution. AIC is derived from asymptotic results (i.e., as sample size approaches infinity; Akaike 1974). With small sample size it tends to be biased in favor of more highly parameterized models. Thus, with limited data, the small sample size version of AIC (AIC$_c$) is recommended (Burnham and Anderson 2002). The formula for AIC$_c$ is given in Eq. (6).

Eq. (6) $\mathrm{AIC}_c = -2L + 2K\left(\frac{n}{n-K-1}\right)$

In Eq. (6), L and K are as above and n is the sample size. The second term on the right-hand side of the above equation is a penalty term. It increases the AIC$_c$ statistic with each additional parameter estimated. Because the denominator of the quotient within the parentheses is zero or negative whenever $n \leq K + 1$, AIC$_c$ cannot be applied to such cases. In practice, n should greatly exceed K when fitting SSDs.

Schwarz (1978) proposed an alternative to AIC that is often referred to as the Bayesian Information Criterion (BIC). It is similar in form and design to AIC and is available in the SSD Toolbox for distributions fit using the Metropolis-Hastings algorithm. The formula for BIC is given in Eq. (7), where K, L, and n are as defined above.

Eq. (7) $BIC = -2L + Kln(n)$

Model averaged HCp

Model-averaged HCp values may be calculated as weighted averages of the HCp values from each individual distribution fit to the same data set using Akaike weights (Δ_i = difference in AIC$_c$ between the i^{th} model and the model with the lowest AIC$_c$, Burnham and Anderson 2002). The formula for Akaike weights is given in Eq. (8).

Eq. (8) $$w_i = \frac{exp\left(-\frac{1}{2}\Delta_i\right)}{\sum_{j=1}^{m} exp\left(-\frac{1}{2}\Delta_j\right)}$$

In the above equation, $\Delta_i = \mathrm{AIC}_c(\text{distribution i}) - min(\mathrm{AIC}_c)$ and the summation is over all (m) distributions compared. Model-averaged estimates of the HCp may be calculated using Eq. (9).

Eq. (9) $\overline{\mathrm{HCp}} = \sum_{j=1}^{m} w_j \mathrm{HCp}_j$

In the above equation, the HCp_j is the estimate of the HCp from the j^{th} distribution considered. Sampling variance of the $\overline{\mathrm{HCp}}$ may be estimated using equation 4.9 of Burnham and Anderson (2002:162), here given as Eq. (10).

Eq. (10) $$var\left(\overline{\mathrm{HCp}}\right) = \sum_{j=1}^{m} w_i \sqrt{var\left(\mathrm{HCp}_j\right) + \left(\mathrm{HCp}_j - \overline{\mathrm{HCp}}\right)^2}$$

Bayesian model-averaging methods are implemented for distributions fit using the Metropolis-Hastings algorithm using the same equations presented above, but with BIC substituted for AIC_c in calculating Δ_i.

Transformations

When the normal, logistic, triangular, or Gumbel distribution are used, the data are first common-log transformed (log_{10}) in the SSD Toolbox. When the Weibull or Burr distribution are used, the data are untransformed. This complicates comparisons among distributions, especially using maximum likelihood and AIC_c. To solve this problem, the likelihoods for the normal, logistic, triangular, and Gumbel distributions are reformulated as follows. First, let:

Eq. (11) $$y = log_{10}(x)$$

Therefore, the cumulative distribution functions for the four distributions using log_{10}-transformed data are of the form: $F(y|\boldsymbol{\theta})$. Thus, the probability density functions for the untransformed data (x) can be calculated using the product rule.

Eq. (12) $$f(x|\boldsymbol{\theta}) = \frac{d}{dx}F(y|\boldsymbol{\theta}) = \frac{d}{dy}F(y|\boldsymbol{\theta})\frac{dy}{dx} = f(y|\boldsymbol{\theta})\frac{1}{x\,ln(10)}$$

In Eq. (12), the expressions $f(y|\boldsymbol{\theta})$ are the probability densities for the log_{10}-transformed data for the respective distributions. When the likelihood is maximized over the transformed data the transformation factor $\frac{1}{x\,ln(10)}$ can be ignored for the purposes of obtaining the MLEs because the transformation factor does not contain the parameters of interest. However, to compare distributions using AIC_c or BIC the factor must be included so that the AIC_c values are on the same scale. Thus, for distributions on log_{10}-transformed data, the log-likelihood on the untransformed scale is given by Eq. (13).

Eq. (13) $$L(\boldsymbol{\theta}|\boldsymbol{X}) \propto \sum_{i=1}^{n} ln\left(\frac{1}{x_i\,ln(10)}f(y|\boldsymbol{\theta})\right) = L(\boldsymbol{\theta}|\boldsymbol{Y}) - ln(10)\sum_{i=1}^{n} y_i - n\,ln(ln(10))$$

Distributions

The SSD toolbox contains functions for fitting six distributions (normal, logistic, triangular, Gumbel, Weibull, and $Burr_{III}$). Table 4 gives some standard statistical notation used in describing the distributions. Table 5 gives a list of distribution functions available in the SSD toolbox.

Table 4. Statistical notation for the description of distributions tested as candidates for use in estimating the HC_p

Symbol	description
n	sample size
$\bar{x}$	sample mean: $\frac{1}{n}\sum_{i=1}^{n} x_i$
s	sample standard deviation: $\sqrt{\frac{1}{n-1}\sum_{i=1}^{N}(x_i - \bar{x})^2}$
$\exp(x)$	exponential function (e^x)
$\mathbf{X}$	column-vector of untransformed data (mean toxicity values)
$\mathbf{Y}$	column-vector of $\log_{10}$-transformed data (mean toxicity values)
$\boldsymbol{\theta}$	column-vector of parameters for any given distribution
$f(x\|\boldsymbol{\theta}), f(y\|\boldsymbol{\theta})$	probability density at x (y if transformed) conditional on $\boldsymbol{\theta}$
$F(x\|\boldsymbol{\theta}), F(y\|\boldsymbol{\theta})$	cumulative distribution function at x (y if transformed) conditional on $\boldsymbol{\theta}$
$F^{-1}(x\|\boldsymbol{\theta}), F^{-1}(y\|\boldsymbol{\theta})$	quantile function at x (y if transformed) conditional on $\boldsymbol{\theta}$
$L(\boldsymbol{\theta}\|\mathbf{X}), L(\boldsymbol{\theta}\|\mathbf{Y})$, or $\boldsymbol{L}$	log-likelihood for $\boldsymbol{\theta}$ conditional on $\mathbf{X}$ ($\mathbf{Y}$ if transformed)

Table 5. Distribution functions in the SSD toolbox

distribution	pdf	cdf	quantile	likelihood	moments	random variates
normal	[1]normpdf	[1]normcdf	[1]norminv	normlik	normmom	[1]randn
logistic	logipdf	logicdf	logiinv	logilik	logimom	logirnd
[2]triangular	triapdf	triacdf	triainv	trialik	trimom	triarnd
Gumbel	gumpdf	gumcdf	guminv	gumlik	gummom	gumrnd
Weibull	[1]wblpdf	[1]wblcdf	[1]wblinv	wbllik	n/a	[1]wblrnd
Burr	burpdf	burcdf	burinv	burlik	n/a	burrnd

[1]functions included in standard Matlab (pdf = probability density function, cdf = cumulative distribution function).

[2]Triangular functions written for the symmetric triangular distribution only

Normal distribution

Parameters: $\boldsymbol{\theta} = [\mu; \sigma]$

μ (location)

σ (scale)

Transformation: $y_i = log_{10}(x_i)$

$$f(y|\mu,\sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} exp\left(-\frac{(y-\mu)^2}{2\sigma^2}\right)$$

$$F(y|\mu,\sigma) = \int_{-\infty}^{y} f(z)\, dz$$

$$F^{-1}(p) = \Phi(p)$$

$$L(\boldsymbol{\theta}|\boldsymbol{Y}) \propto -\frac{n}{2} ln(2\pi) - n\, ln(\sigma) + \frac{1}{2\sigma^2}\sum_{i=1}^{n}(y_i-\mu)^2$$

$$\frac{dL}{d\mu} = -\frac{1}{\sigma^2}\sum_{i=1}^{n}(y_i-\mu)$$

$$\frac{dL}{d\sigma} = -\frac{1}{\sigma}\left(n + \frac{1}{\sigma^2}\sum_{i=1}^{n}(y_i-\mu)^2\right)$$

Mean = μ and Variance = σ^2

Neither the cdf (F) nor the quantile function (F^{-1}) has explicit form. However, both can be readily approximated to arbitrary precision in most mathematical software.

Linearization for graphical estimation makes use of the z-scores, which are the percentiles of a standard normal distribution with mean 0 and unit variance. These are given by the equation $z = \frac{y-\mu}{\sigma}$, which yields the linear equation $y = \sigma z + \mu$. Given z and y, σ and μ can be estimated using linear regression.

Moment Estimators:

$$\hat{\mu} = \bar{y} \qquad \hat{\sigma} = s$$

Logistic distribution

Parameters: $\boldsymbol{\theta} = [\alpha; \beta]$

α (location)

β (scale)

Transformation: $y_i = log_{10}(x_i)$

$$f(y: \alpha, \beta) = \frac{exp(-(y-\alpha)/\beta)}{\beta(1+exp(-(y-\alpha)/\beta))^2}$$

$$F(y: \alpha, \beta) = \frac{1}{1+exp(-(y-\alpha)/\beta)}$$

$$F^{-1}(p) = \alpha + \beta\, ln\left(\frac{p}{1-p}\right)$$

Let:

$$r_i = y_i - \alpha \quad \text{and} \quad m_i = exp\left(-\frac{r_i}{\beta}\right)$$

$$L(\boldsymbol{\theta}|\boldsymbol{Y}) \propto \frac{1}{\beta}\sum_{i=1}^{n}(r_i) - n\, ln(\beta) - 2\sum_{i=1}^{n} ln(1+m_i)$$

$$\frac{dL}{d\alpha} = \frac{n}{\beta} - \frac{2}{\beta}\sum_{i=1}^{n}\left(\frac{m_i}{1+m_i}\right)$$

$$\frac{dL}{d\beta} = \frac{1}{\beta^2}\sum_{i=1}^{n} r_i - \frac{n}{\beta} - \frac{2}{\beta^2}\sum_{i=1}^{n}\frac{r_i m_i}{1+m_i}$$

Mean = α and Variance = $\frac{\pi^2}{3}\beta^2$

Linearization for graphical estimation can be done using a standard logistic distribution ($\alpha = 0$, $\beta = \frac{\sqrt{3}}{\pi}$), with standard quantiles (z_L) defined as $z_L = \sqrt{3}\frac{y-\alpha}{\pi\beta}$, which yields the linear equation $y = \frac{\pi}{\sqrt{3}}\beta z_L + \alpha$.

Moment Estimators:

$$\hat{\alpha} = \bar{y} \text{ and } \quad \hat{\beta} = s\frac{\sqrt{3}}{\pi}$$

Triangular distribution (symmetric)

Parameters: $\boldsymbol{\theta} = [a\,;b]$

a (minimum)

b (maximum)

Transformation: $y_i = log_{10}(x_i)$

If: $a \le y \le \frac{a+b}{2}$:

$$f(y|a,b) = \frac{4(y-a)}{(b-a)^2}$$

$$F(y|a,b) = \frac{2(y-a)^2}{(b-a)^2}$$

$$L(\boldsymbol{\theta}|\boldsymbol{Y}) \propto -2\,ln(b-a) + ln(4) + \sum_{i=1}^{n}(y_i - a)$$

If: $\frac{a+b}{2} < y \le b$:

$$f(y|a,b) = \frac{4(b-y)}{(b-a)^2}$$

$$F(y|a,b) = 1 - \frac{2(y-b)^2}{(b-a)^2}$$

$$L(\boldsymbol{\theta}|\boldsymbol{Y}) \propto -2\,ln(b-a) + ln(4) + \sum_{i=1}^{n}(b - y_i)$$

If: $p \le 0.5$

$$F^{-1}(p) = a + \sqrt{\frac{p(b-a)^2}{2}}$$

If: $p > 0.5$

$$F^{-1}(p) = b + \sqrt{\frac{(1-p)(b-a)^2}{2}}$$

Mean = $\frac{a+b}{2}$ and Variance = $\frac{(b-a)^2}{24}$

Linearization of the triangular distribution makes use of the standard symmetric triangular distribution ($a = -\sqrt{6}$ and $b = \sqrt{6}$). Defining percentiles of the standard symmetric triangular as z_T, we have $z_T = \frac{y - \frac{a+b}{2}}{\frac{b-a}{\sqrt{24}}}$, which yields the linear equation: $y = z_T \frac{(b-a)}{\sqrt{24}} + \frac{(a+b)}{2}$. This is a linear equation with $slope = \frac{(b-a)}{\sqrt{24}}$ and $intercept = \frac{(a+b)}{2}$, which can be estimated using linear regression. Linear substitution can then be used to solve for a and b:

$$\hat{a} = intercept - slope\sqrt{6} \quad \text{and} \quad \hat{b} = intercept + slope\sqrt{6}$$

Moment Estimators:

$$\hat{a} = \bar{y} - s\sqrt{6} \quad \text{and} \quad \hat{b} = \bar{y} + s\sqrt{6}$$

Gumbel (Gompertz, Extreme Value Type 1) distribution

Parameters: $\boldsymbol{\theta} = [\mu, \beta]$

μ (location)

β (scale)

Transformation: $y_i = log_{10}(x_i)$

$$F(y) = exp\left(-exp\left(\frac{\mu - y}{\beta}\right)\right)$$

$$f(y) = \frac{1}{\beta} exp\left(\frac{\mu - y}{\beta} - exp\left(\frac{\mu - y}{\beta}\right)\right)$$

$$F^{-1}(p) = \mu - \beta\, ln(-ln(p))$$

Let:

$$z_i = \frac{\mu - y_i}{\beta}$$

Then:

$$L(\boldsymbol{\theta}|\boldsymbol{Y}) \propto -n\, ln(\beta) + \sum_{i=1}^{n} z_i - \sum_{i=1}^{n} exp(z_i)$$

Note:

$$\frac{dL}{dz_i} = -1 + exp(z_i) \qquad \frac{dz_i}{d\mu} = \frac{1}{\beta} \qquad \text{and} \qquad \frac{dz_i}{d\beta} = -\frac{\mu - y_i}{\beta^2}$$

$$\frac{dL}{d\mu} = \frac{1}{\beta}\sum_{i=1}^{n}(1 - exp(z_i)) \qquad \text{and} \qquad \frac{dL}{d\beta} = -\frac{n-1}{\beta}\sum_{i=1}^{n} z_i(1 - exp(z_i))$$

Mean:

$\mu + \beta\gamma$ where γ = Euler-Mascheroni constant.

Variance

$$\frac{\beta^2\pi^2}{6}$$

Linearization of the Gumbel quantile function can be performed directly by setting $F(y) = ECDF(y)$, which yields the linear equation $y = \beta(-ln[-ln(ECDF(y)]) + \mu$.

Moment Estimators:

$$\hat{\beta} = \frac{s}{\pi}\sqrt{6} \qquad \text{and} \qquad \hat{\mu} = \bar{y} - \hat{\beta}\gamma$$

Weibull (Extreme Value Type III) distribution

Parameters: $\boldsymbol{\theta} = [\lambda; k]$

λ (scale)

k (shape)

Transformation: none

$$\boldsymbol{F(x) = 1 - exp\left(-\left(\frac{x}{\lambda}\right)^{k}\right)}$$

$$\boldsymbol{f(x) = \frac{k}{\lambda}\left(\frac{x}{\lambda}\right)^{k-1} exp\left(-\left(\frac{x}{\lambda}\right)^{k}\right)}$$

$$\boldsymbol{F^{-1}(p) = \lambda(-ln(1-p))^{1/k}}$$

$$\boldsymbol{L(\theta|X) = n\,ln(k) - kn\,ln(\lambda) + (k-1)\sum_{i=1}^{n} ln(x_i) - \sum_{i=1}^{n}\left(\frac{x_i}{\lambda}\right)^{k}}$$

$$\boldsymbol{\frac{dL}{d\lambda} = -\frac{k}{\lambda}\left(n - \sum_{i=1}^{n}\left(\frac{x_i}{\lambda}\right)^{k}\right)}$$

$$\boldsymbol{\frac{dL}{dk} = \frac{n}{k} - n\,ln(\lambda) + \sum_{i=1}^{n} ln(x_i) - \sum_{i=1}^{n}\left(\left(\frac{x_i}{\lambda}\right)^{k} ln\left(\frac{x_i}{\lambda}\right)\right)}$$

The mean and variance of the Weibull distribution are:

$$\boldsymbol{\lambda\Gamma\left(1+\frac{1}{k}\right)} \quad \text{and} \quad \boldsymbol{\lambda^2\Gamma\left(1+\frac{2}{k}\right) - \left(\lambda\Gamma\left(1+\frac{1}{k}\right)\right)^2}$$

In the above equations Γ is the gamma function.

Linearization of the Weibull distribution is accomplished by setting $F(x) = ECDF(x)$, which yields the linear equation $ln(x) = \frac{1}{k} ln\left(-ln\left(\mathbf{1} - ECD\boldsymbol{F(x)}\right)\right) + ln(\lambda)$.

The gamma functions in the equations for the mean and variance prevent moment estimators from being derived.

$Burr_{III}$ distribution

Note: this is the $Burr_{III}$ distribution from Shao (2000)

Parameters: $\boldsymbol{\theta} = [b\,;c\,;k]$

$$F(x) = \frac{1}{\left[1+\left(\frac{b}{x}\right)^{c}\right]^{k}}$$

$$f(x) = \frac{kc}{b}\frac{\left(\frac{b}{x}\right)^{c+1}}{\left[1+\left(\frac{b}{x}\right)^{c}\right]^{k+1}}.$$

$$\boldsymbol{F^{-1}}(\boldsymbol{p}) = b\left(p^{-1/k}-1\right)^{-1/c}$$

Note that Shao (2000, Eq. 8) incorrectly gives the pdf as: $f(x) = \frac{kc}{b}\frac{\left(\frac{b}{x}\right)^{c+1}}{\left[1+\left(\frac{b}{x}\right)^{c+1}\right]^{k+1}}$

$$L(\boldsymbol{\theta}|\boldsymbol{X}) \propto n\,ln(c) + n\,ln(k) + cn\,ln(b) - (c+1)\sum_{i=1}^{n} ln(x_i) - (k+1)\sum_{i=1}^{n} ln\left(1+\left(\frac{b}{x_i}\right)^{c}\right)$$

Let:

$$z_i = 1+\left(\frac{b}{x_i}\right)^{c}$$

Then:

$\frac{dz_i}{db} = \frac{c}{b}\left(\frac{b}{x_i}\right)^{c}$ and

$\frac{dz_i}{dc} = ln\left(\frac{b}{x_i}\right)\left(\frac{b}{x_i}\right)^{c}$

$\frac{dL}{db} = \frac{cn}{b} - (k+1)\sum_{i=1}^{n}\frac{1}{z_i}\frac{dz_i}{db}$

$$\frac{dL}{dc} = \frac{n}{c} + n\,ln(b) - \sum_{i=1}^{n} ln(x_i) - (k+1)\sum_{i=1}^{n}\frac{1}{z_i}\frac{dz_i}{dc}$$

$$\frac{dL}{dk} = \frac{n}{k} - \sum_{i=1}^{n} ln(z_i)$$

Moment estimators are not available for the $Burr_{III}$ distribution.
Linearization methods are not available for the $Burr_{III}$ distribution.

Literature cited

Akaike, H. 1974. A new look at the statistical model identification. IEEE Transactions on Automatic Control 19: 716–723.

Aldenberg, T. and R. Luttik. 2002. Extrapolation Factors for Tiny Toxicity Data Sets from Species Sensitivity Distributions with Known Standard Deviation. Pp. 103-117 in L. Posthuma, G.W. Suter, and T.P. Traas (Eds.) Species Sensitivity Distributions in Ecotoxicology. Lewis. Boca Raton, FL. USA.

Aldenberg, T. and W. Slob. 1993. Confidence limits for hazardous concentrations based on logistically distributed NOEC toxicity data. Ecotoxicology and Environmental Safety 25:48-63.

Arnold, B.C., N. Balakrishnan, and H.H. Nagaraja. 2008. A first course in order statistics. Society for Industrial and Applied Mathematics. Philadelphia, PA, USA.

Burnham, K.P. and D.R. Anderson. 2002. Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach, 2nd Ed. Springer Verlag, New York, NY.

Campbell, E., M. Palmer, Q. Shao, and D. Wilson. 2000. BurrliOZ. CSIRO, Mathematical and Information Sciences. www.cmis.csiro.au/envir/burrlioz.

Chapman, P.F., M. Reed, A. Hart, W. Roelofs, T. Aldenberg, K. Solomon, J. Tarazona, M. Liess, P. Byrne, W. Powley, J. Green, S. Ferson, and H. Galicia. 2007. Methods of Uncertainty Analysis. Work Package 4 in EUFRAM: Concerted action to develop a European framework for probabilistic risk assessment of the environmental impacts of pesticides. Version 3, March 2007. http://www.eufram.com/, accessed 8 August 2011.

Edwards, A.W.F. 1992. Likelihood: Expanded edition. Johns Hopkins University Press. Baltimore, MD, USA.

Efron, B. and R.J. Tibshirani. 1994. An Introduction to the Bootstrap. Chapman and Hall/CRC. Boca Raton, FL. USA.

Elphick, J.R.F., K.D. Bergh, and H.C. Bailey. 2011. Chronic toxicity of chloride to freshwater species: effects of hardness and implications for water quality guidelines. Environmental Toxicology and Chemistry 30:239-246.

Erickson, R.J. and C.E. Stephan. 1988. Calculation of the final acute value for water quality criteria for aquatic organisms. U.S. Environmental Protection Agency. Duluth, MN. EPA/600/3-88-018.

Fojut, T.L, A.J. Palumbo, and R.S. Tjeerdema. 2012. Aquatic life water quality criteria derived via the UC Davis method: II. Pyrethroid Insecticides. Reviews of Environmental Contamination and Toxicology 2016:51-104.

Harville, D.A. 1977. Maximum likelihood approaches to variance component estimation and to related problems. Journal of the American Statistical Association 72:320-340.

Hastings, W.K. 1970. Monte Carlo Sampling Methods Using Markov Chains and Their Applications. Biometrika. 57:97-109.

Kooijman, S.A.L.M. 1987. A safety factor for LC_{50} values allowing for differences in sensitivity among species. Water Research 21:269-276.

King, R., B.J.T. Morgan, O. Gimenez, and S. P. Brooks. 2010. Bayesian Analysis for Population Ecology. CRC Press, Boca Raton, FL.

Link, W.A. and R.J. Barker. 2010. Bayesian Inference with Ecological Applications. Elsevier, Academic Press, London, UK.

Luttik, R. and T. Aldenberg. 1997. Extrapolation factors for small samples of pesticide toxicity data: special focus on LD50 values for birds and mammals. Environmental Toxicology and Chemistry 16:1785-1788.

Manly, B.F.J. 1997. Randomization, Bootstrap, and Monte Carlo Methods in Biology, 2nd Ed. Chapman & Hall/CRC, Boca Raton, FL.

Mineau, P., Collins, B.T., and A. Baril. 1996. On the use of scaling factors to improve interspecies extrapolation of acute toxicity in birds. Regulatory Toxicology and Pharmacology, 24: 24-29.

Moore, D.R.J., C.D. Priest, N. Galic, R.A. Brain and S.I. Rodney. 2019. Correcting for phylogenetic autocorrelation in species sensitivity distributions. Integrated Environmental Assessment and Management (in press). DOI:10.1002/ieam.4207.

Newman, M.C., D.R. Ownby, L.C.A. Mézin, D.C. Powell, T.R.L Christensen, S.B. Lerberg, and B.-A. Anderson. 2000. Applying species sensitivity distributions in ecological risk assessment: assumptions of distribution type and sufficient numbers of species. Environmental Toxicology and Chemistry 19:508-515.

Posthuma, L, T.P. Traas, and G.W. Suter. 2002. General introduction to species sensitivity distributions. Pp. 3-10 *in* L. Posthuma, G.W. Suter, and T.P. Traas (Eds.) Species Sensitivity Distributions in Ecotoxicology. Lewis. Boca Raton, FL. USA.

Schwarz, G.E. 1978. Estimating the dimension of a model. Annals of Statistics 6:461–464.

Seber, G.A.F. 2002. The estimation of animal abundance and related parameters, 2nd Ed. Reprint of 1982 edition. Blackburn Press. Caldwell, NJ. USA.

Shao, Q. 2000. Estimation for hazardous concentrations based on NOEC toxicity data: an alternative approach. Environmetrics 11:583-595.

Suter, G.W. 2002. North American history of species sensitivity distributions. Pp. 11-18 *in* L. Posthuma, G.W. Suter, and T.P. Traas (Eds.) Species Sensitivity Distributions in Ecotoxicology. Lewis. Boca Raton, FL. USA.

TenBrook, P. L., R.S. Tjeerdema, P. Hann, and J. Karkoski. 2008. Reviews of Environmental Contamination and Toxicology 199, DOI: 10.1007/978-0-387-09808-1.

TenBrook PL, Palumbo AJ, Fojut TL, Hann P, Karkoski J, Tjeerdema RS. 2010. The University of California-Davis Methodology for Deriving Aquatic Life Pesticide Water Quality Criteria. Reviews of Environmental Contamination and Toxicology. Volume 209.

[USEPA] US Environmental Protection Agency. 2016. Biological evaluation chapters for Chlorpyrifos ESA assessment. [cited 2019 July 3]. https://www.epa.gov/endangered-species/biological-evaluation-chapters-Chlorpyrifos-esa-assessment

van Straalen, N.M and C.A.J. Denneman. 1989. Ecotoxicological evaluation of soil quality criteria. Ecotoxicology and Environmental Safety 18:241-251.

Zajdlik & Associates. 2005. Statistical analysis of the SSD approach for development of Canadian Water Quality Guidelines. Canadian Council of Ministers of the Environment. Project #354-2005.