Estimating ten Berge Exponents

Final Report

Takumi Kijima & Matthew Snyder

EXECUTIVE SUMMARY

High concentrations of dangerous chemicals can cause severe harm or death to humans when exposed. Immediately Dangerous to Life and Health (IDLH) values are chemical-specific measures of the highest concentration of a chemical that a human can survive for 30 minutes. This report uses time scaling to determine the ten Berge exponents, an intermediary metric in computing IDLH values, for eight chemicals based on lethality data collected on rats. These exponents are computed as a function of linear regression parameter estimates and typically range from 0.85 to 3.5, so credible intervals were determined using a Bayesian approach. Six of the eight chemicals (Acrolein, Ammonia, Carbon disulfide, Carbon monoxide, Epichlorohydrin, and Ethyleneimine) had lower bounds below 0.85, while the remaining two chemicals (Hydrogen cyanide and Pentaborane) had credible intervals entirely contained within the expected range.

INTRODUCTION

Immediately Dangerous to Life and Health (IDLH) values for different chemicals represent the highest concentration level that is survivable for humans for at least 30 minutes. These values are determined from lethality data collected on animals from different exposure times. This project focuses on the first step of the process, in which the lethality limits are converted to a 30-minute exposure time using time-scaling methods. Two methods are available: Haber's Rule and ten Berge. Haber's Rule states that the incidence or severity of a health outcome equals $C \times t$, where C is the concentration and t is the time, while ten Berge's rule states that it equals $C^n \times t$, where n is the ten Berge exponent. The value of n is chemical specific but is typically assumed to be constant across species for a given chemical; values are expected to range between 0.85 and 3.5, and a value of 1 reduces to Haber's Rule. The primary question of interest for this project is to determine the ten Berge exponent for different chemicals using linear regression.

There were 30 different chemicals provided in the data. Each observation corresponds to an experiment and contains the study reference, species, chemical, and the time and LC_{50} obtained from the experiment. This analysis focuses on data collected on rats, but other species were considered for validation purposes. All observations with missing LC_{50} and time values were

removed, as well as all observations for chemicals in which all experiments were performed for the same time exposure, since time scaling cannot be performed. After this data-cleaning procedure, only 8 of the 30 chemicals had at least three observations from three unique exposure times, the minimum amount of data necessary to obtain informative results. These chemicals are summarized in Table 1 below.

Table 1. Summary of Data by Chemical.

Chemical	Species	Number of Observations	Number of Unique Time Exposures
Acrolein	Rat	3	3
Ammonia	Rat	6	5
Carbon disulfide	Rat	4	3
Carbon monoxide	Rat	3	3
Epichlorohydrin	Rat	5	3
Ethyleneimine	Rat	3	3
Hydrogen cyanide	Rat	5	3
Pentaborane	Rat	5	5
	Mouse	6	5

METHODOLOGY

Linear Regression

Based on the Standard Operating Procedures and the exploratory analysis seen in Figure 1 of the Results, the relationship between the log of the time and the log of the LC_{50} values appear linear, so the data was transformed for analysis. By taking the natural log of the ten Berge equation, we have $log(C^n \times t) = log(k)$, for a constant k, which implies that $log(C) = \frac{log(k)}{n} - \frac{1}{n} log(t)$ by the rules of logarithms. Then, since the relationship between $log(C)$ and $log(t)$ is linear, the process can be modeled using simple linear regression for each chemical and each species. Thus, we will fit the following model:

$$log(C_i) = \beta_0 + \beta_1 log(t_i) + \epsilon_i,$$

where i is the number of observations for each chemical on the species, C_i is the LC_{50} concentration in parts per million (ppm), t_i is the time in minutes, β_0 is the intercept, and β_1 represents the

change in the log concentration for each log-minute increase in exposure time. Based on the ten Berge equation, β_1 equals $-\frac{1}{n}$. By assumption, $\epsilon_i \sim N(0, \sigma^2)$.

To estimate β_1, and then n, a Bayesian approach was used. Bayesian analysis is a statistical technique that combines prior knowledge with information from the data in order to obtain an updated distribution for our value, called the posterior distribution. Since it is known that n is typically between 0.85 and 3.5, β_1 is expected to be between -1.1765 and -0.2857. Thus, we assume that the prior distribution of β_1 is normal, with a mean of -0.7311 and standard deviation is 0.2272. The mean and standard deviation were determined so that β_1 is between -1.1765 and -0.2857 with a probability of 0.95; this implies that 95 percent of the time, the ten Berge exponent is between 0.85 and 3.5. To obtain the posterior distribution, Markov Chain Monte Carlo (MCMC) was used. An estimate of β_1, $\widehat{\beta_1}$, was obtained by taking the mean of this posterior distribution, and n was then estimated by taking $-\frac{1}{\widehat{\beta_1}}$.

Linear models assume linearity and independence in data collection, and constant variance and normality in the residuals. The linearity assumption is generally accepted in this research area and will be checked visually. Independence was assumed, although there may be violations due to multiple measurements taken during the same study. Constant variance and normality were checked using the Residuals vs. Fitted plot and Normal Q-Q plots from the regression output.

Measures of Uncertainty

Since the ten Berge exponent is computed as a function of β_1, several different techniques were used and compared in order to measure the uncertainty of the exponent estimates. These approaches, including jackknifing, bootstrapping, and Bayesian methods, are described and compared in Appendix A. Based on a simulation study, the High Density Bayesian Interval (HDI) method is recommended due to its high accuracy and small interval widths, and will be used in this report.

The HDI credible intervals were determined from the posterior distribution obtained using the Bayesian approach, described above. This approach finds the interval of the distribution that contains 95 percent of the density, while also ensuring that each point in the interval has a higher density than any point outside of the interval. For skewed distributions, this method results in

narrower intervals than traditional quantile approaches, while still containing the true value with 95 percent probability.

Model Validation

From a biological perspective, ten Berge exponents are expected to be constant across species for a given chemical. Thus, for chemicals having sufficient data on multiple species, the models were fit using one species and the intervals for the ten Berge exponents were computed. Then, the same model was fit to data from the same chemical but on other available species, and the corresponding ten Berge exponent was compared to the interval from the original model. If the new estimate falls within the original interval, the model structure is deemed to be appropriate.

RESULTS

Exploratory Data Analysis

For the eight chemicals with sufficient data, the LC_{50} values were plotted with respect to time; without transforming the data, no patterns emerged. Based on the standard practice of transforming the time and LC_{50} values using natural log, we recreated the plots by chemical, which are shown in Figure 1. These show that the relationship between the log of the LC_{50} values in parts per million (ppm) and the log of the time in minutes is relatively linear for all chemicals, although the rates of change vary.

Figure 1: LC_{50} Values for Rats over Time by Chemical.

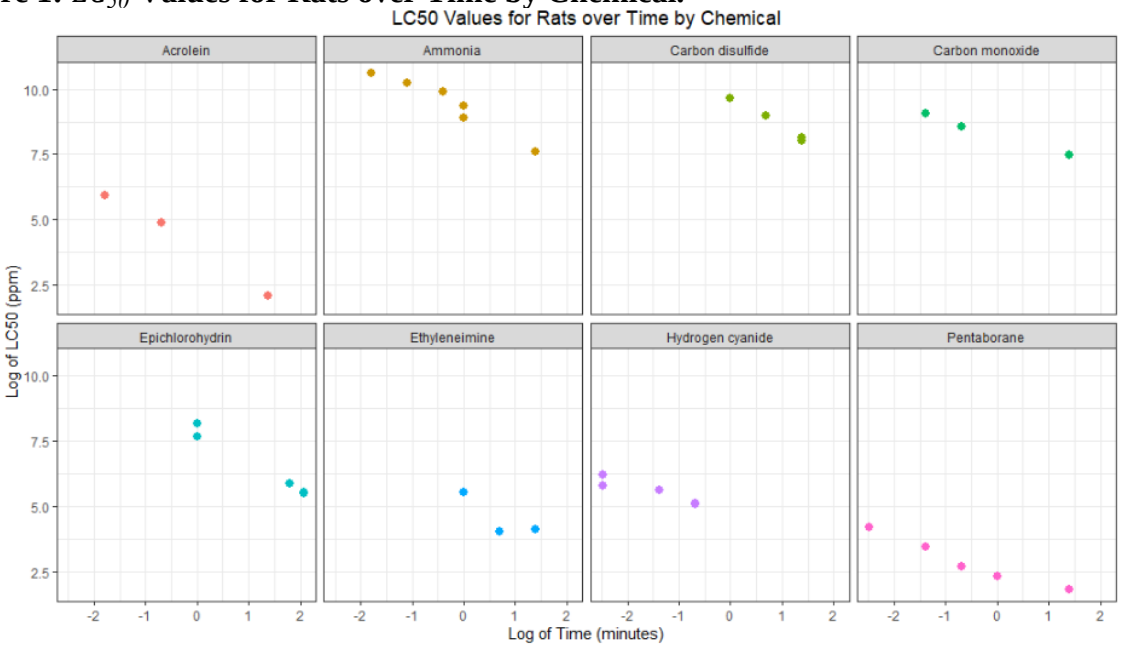

In addition to the above plots for the rat experiments, Pentaborane was the only chemical that had at least three observations for another species. Figure 2 compares the LC_{50} values over time between the experiments on mice, in red, and rats, in blue. While the values for rats are uniformly larger, the slopes of the lines appear fairly similar between the species; this confirms the belief that the relationship between the time and lethality concentration varies across chemicals but is consistent across species.

Figure 2: LC_{50} Values for Mice and Rats over Time for Pentaborane.

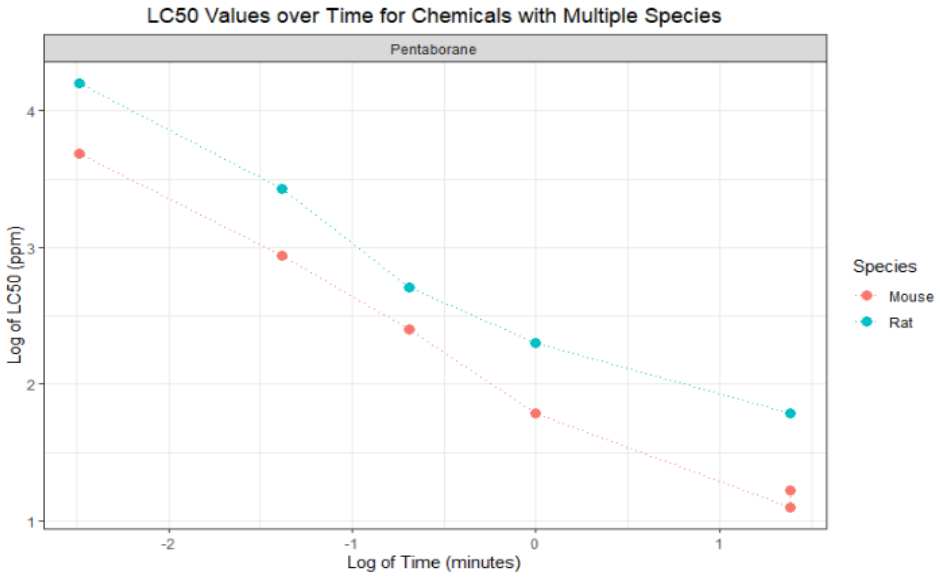

Linear Regression

Based on the relationships seen in Figure 1 and the structure of the formula described in the Methodology, a log-linear regression model was fit for each chemical with the log LC_{50}values as a function of the log of the time. Residual vs. Fitted and Normal Q-Q plots are provided in Appendix B to check the assumptions. The fitted models are shown for each of the eight chemicals in Figure 3 below. The values of $\widehat{\beta_0}$ and $\widehat{\beta_1}$ used in the fitted models were determined from their respective posterior distributions from the Bayesian analysis.

Figure 3: Fitted LC_{50} Values for Rats over Time by Chemical

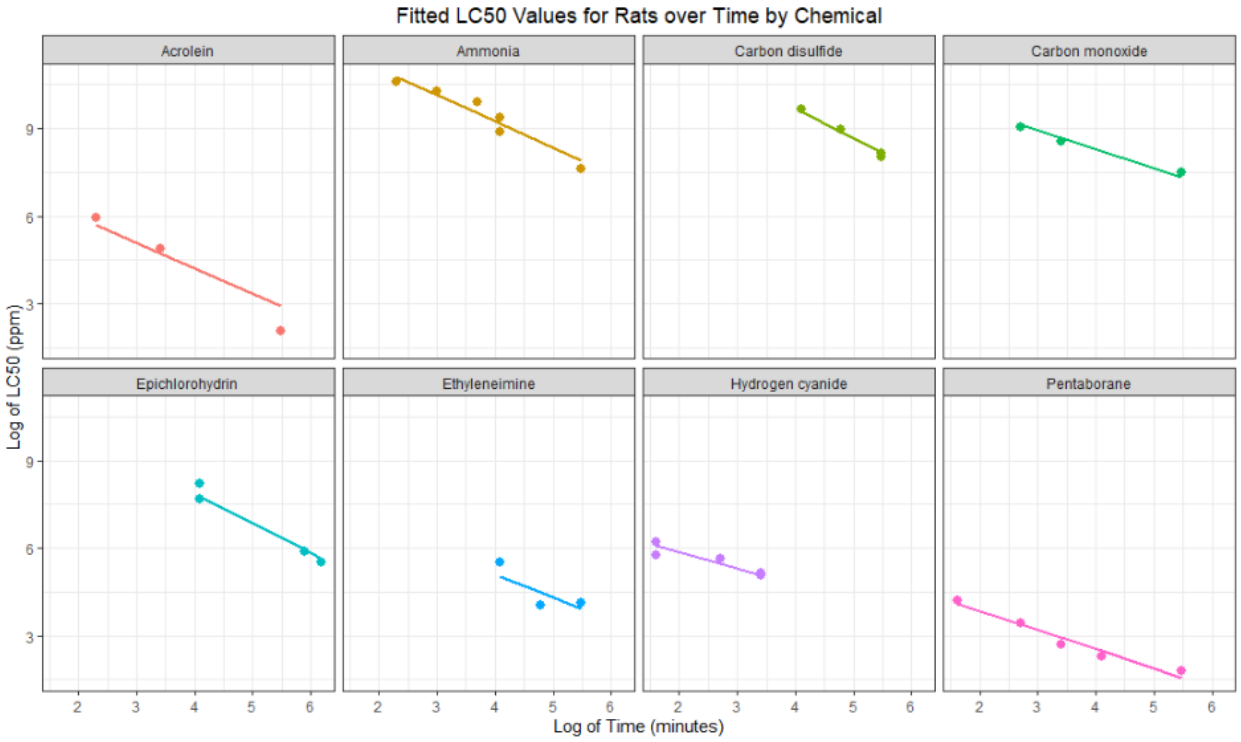

Measures of Uncertainty

The five methods for determining intervals were compared using simulation and the results are described in Appendix A. Although the Bayesian credible intervals were often slightly wider than the confidence intervals, they were similar in magnitude and less prone to unreasonable bounds caused by anomalies in the small datasets. In addition, the HDI Bayesian method had the highest coverage probabilities and smallest average widths for the sample sizes considered, in general; thus, this method was used for the analysis. Based on this method, Table 2 below shows the number of observations, the β_1 coefficient estimate from the linear model, the computed ten Berge exponent (n), and the 95 percent credible interval for the ten Berge exponent for each of the eight chemicals.

Table 2: Log-Linear Regression Output and ten Berge Exponent Estimation.

Chemical	Number of Observations	$\widehat{\beta_1}$	*n*	Credible Interval for *n*
Acrolein	3	-1.23	1.18	(0.74, 2.17)
Ammonia	6	-0.97	1.13	(0.83, 1.58)
Carbon disulfide	4	-1.14	1.03	(0.71, 1.61)
Carbon monoxide	3	-0.56	1.46	(0.81, 2.41)
Epichlorohydrin	5	-1.16	1.01	(0.75, 1.51)
Ethyleneimine	3	-1.01	1.31	(0.75, 2.53)
Hydrogen cyanide	5	-0.49	1.73	(0.98, 2.88)
Pentaborane	5	-0.64	1.51	(1.06, 2.13)

Further, Figure 4 displays the credible interval estimates for the ten Berge exponents for each of the eight chemicals, along with their point estimates. The solid horizontal lines occur at values of 0.85 and 3.5, framing the generally accepted range, and the dotted lines occur at 1 and 3, marking the boundaries of the frequently used range for the exponents.

Figure 4: Estimates and 95% Credible Intervals for ten Berge Exponents.

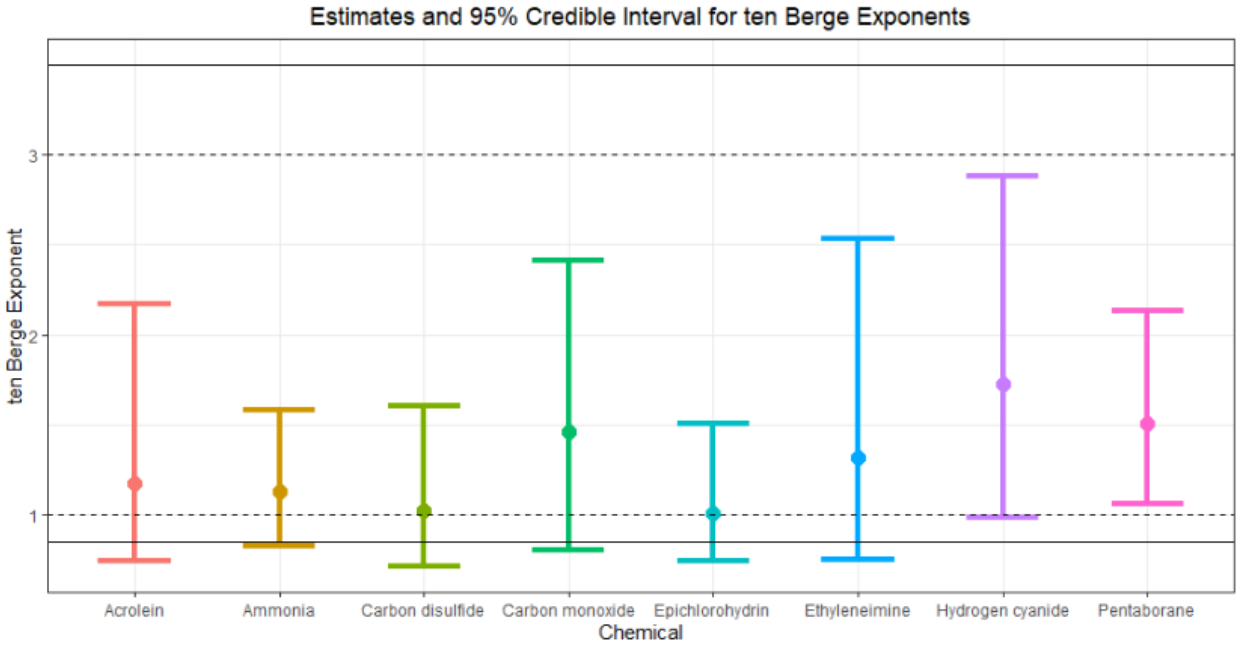

Model Validation

Pentaborane was the only chemical with sufficient data to perform the analysis on a species other than rats, so the same log-linear regression model was fit on the mouse data. This resulted in a β_1 estimate of -0.66, for an estimated ten Berge exponent of 1.51. Recall that the estimated exponent based on the rat data was 1.51, with a credible interval of (1.06, 2.13). Thus, the estimated exponent based on the mouse data falls within the credible interval, so the model structure appears to be appropriate for this analysis, since it provides similar results for data collected on different species for the same chemical.

DISCUSSION

Acrolein, Carbon monoxide, and Ethyleneimine only have three valid observations, while the other five chemicals have at least four. In each case, there is a negative relationship between the log of the time and the log of the LC_{50} values, although the magnitude of the slope varies. This implies positive ten Berge exponents, which are shown in Table 2, which range from 1.01 to 1.73. Further, the HDI Bayesian credible intervals for six of the eight chemicals (Acrolein, Ammonia, Carbon disulfide, Carbon monoxide, Epichlorohydrin, and Ethyleneimine) contain values below 0.85 but above 0.70. Intervals for Hydrogen cyanide and Pentaborane fall within the generally accepted 0.85 to 3.5 range, and Pentaborane is the only chemical with a credible interval entirely contained within 1 and 3.

Of the four widest credible intervals, three correspond to the three chemicals with only three observations (Acrolein, Carbon monoxide, and Ethyleneimine). Further, the smallest credible interval in this set of chemicals belongs to Ammonia, which had the most data, at six observations and five unique time points. This indicates that as the number of observations increases, the width of the 95 percent credible intervals decreases, providing more insight to the true ten Berge exponent; from the Simulation study in Appendix A, these credible intervals are still expected to contain the true value over 94 percent of the time when three to ten observations are available.

In terms of model validation, since the estimated ten Berge coefficient for Pentaborane using the mouse data was contained within the credible interval generated based on the rat data, this modeling strategy seems reasonable and provides consistent results. However, obtaining more data from other species and other chemicals is necessary to generalize this claim.

An important assumption in linear regression is the independence of observations. However, since some observations can come from the same study, this assumption is not always satisfied. Future work that accounts for the dependence by using random effects in the regression model could be performed in order to satisfy this assumption.

An additional concern in this analysis is the small number of observations for each chemical. In particular, the assumptions of normality and constant variance needed for linear regression cannot be thoroughly analyzed with small sample sizes; although there are no obvious violations seen in the plots in Appendix B, this analysis relies heavily on these assumptions. Further, the fewer the observations, the wider the credible intervals. Thus, more observations would allow for narrower and more accurate interval estimates for the ten Berge exponents.

REFERENCES

The regression modeling is based on the linear regression techniques described in *Applied Linear Statistical Models, 4th Edition* by Kutner, et al, and the Bayesian method is based on *Bayesian Data Analysis, 4th Edition* by Gelman et al. The bootstrapped confidence interval approaches are based on the article "Understanding Bootstrap Confidence Interval Output from the R boot Package" by Jeremy Albright[1].

All data handling, visualization, and analysis were performed using R (64-bit, version 4.0.0). The 'tidyverse' package (version 1.3.0) was used for data manipulation and plotting, and the 'boot' package (version 1.3-25) was used to construct the bootstrapped confidence intervals. The 'rstan' (version 2.21.2), 'bayesplot' (version 1.7.2), and 'HDInterval' (version 0.2.2) packages were used for the Bayesian analysis and construction of credible intervals.

[1] "Understanding Bootstrap Confidence Interval Output." Methods. 13 Sept. 2019, blog.methodsconsultants.com/posts/understanding-bootstrap-confidence-interval-output-from-the-r-boot-package/.

APPENDIX A: Simulation Study

In this analysis, intervals for the ten Berge exponents were desired to account for the variability in the estimates. However, since these exponents are computed as a function of the β_1 estimate from linear regression and the sample sizes are very small, traditional confidence intervals are not appropriate. Alternative approaches include empirically based confidence intervals obtained from resampling methods such as bootstrapping or jackknifing. In addition, since the typical range for the exponents is known, a Bayesian approach can produce credible intervals that account for this prior knowledge. Since there is not a consensus in the literature about the best approach for obtaining intervals on small data, a simulation study was conducted. Five methods – jackknife, percentile bootstrap, bias-corrected and accelerated (BCa) bootstrap, ETI Bayesian, and HDI Bayesian – were compared at various sample sizes and the coverage probabilities were computed.

Methodology

These methods include confidence intervals derived from resampling techniques such as jackknifing and bootstrapping, as well as credible intervals obtained from a Bayesian approach that accounts for prior knowledge about the exponents. A 95 percent confidence interval is constructed so that, out of 100 intervals, 95 contain the true value. On the other hand, a 95 percent credible interval is constructed so that the true value is contained in the interval with 95 percent probability. Although fundamentally different, these two methods are used to measure the uncertainty in the parameter estimates.

Jackknifing is a process of removing one observation from the data, computing the desired statistic, and repeating this process to obtain a distribution of possible values. A 95 percent confidence interval can be obtained by determining the 2.5 and 97.5 quantiles from the distribution. For this analysis, one observation was removed from the data, the regression was performed, and n was computed as a function of the β_1 estimate. This process was then repeated to obtain the full distribution of exponent values and the confidence interval was computed.

A similar technique is bootstrapping, in which the observations are randomly sampled from the data with replacement to form new datasets of equal size. The regression model was fit on these samples, and the ten Berge exponents were calculated. Again, this process was repeated to obtain

an exhaustive list of values, and a 95 percent confidence interval was determined from the 2.5 and 97.5 quantiles. This is referred to as the percentile method of the bootstrapping technique.

A third approach is an extension of the percentile bootstrap method, called the Bias-Corrected and Accelerated bootstrap method, or BCa. In addition to the traditional bootstrap process, BCa confidence intervals account for both the bias and skewness of the resampled distribution by accounting for the original estimate from the data. The exact bias correction and acceleration terms are described in *Understanding Bootstrap Confidence Interval Output*[2]. For larger samples, the BCa approach is typically preferred[3] due to the additional corrections; however, there is no consensus in the literature for the correct method when the sample size is this small.

The final two approaches are based on the Bayesian method: HDI, as described in the Methodology of the main report, and Equal Tail Intervals (ETI). Using the posterior distribution from the Bayesian analysis as described in the Methodology section of the main report, the ETI method computes credible intervals in a similar manner as the percentile bootstrap method. The 2.5 and 97.5 quantiles of the posterior distribution are used as the lower and upper bounds of the 95 percent credible interval, so each tail has the same probability. For a symmetric distribution, the ETI and HDI methods are equivalent; however, when the distribution is skewed, the HDI method results in slightly narrower intervals, since the tails need not be equal.

The five approaches were compared by their coverage probability, which is the percentage of the simulated intervals that contain the true value out of 1,000 samples, and their average width. Higher coverage probabilities indicate more accurate interval estimates, while smaller widths provide more insight into the true value of the parameter. This process was repeated for several true values of *n*, and the method with the highest coverage probability and smallest interval width, in general, is recommended and used in this analysis.

As a basis for the simulation, the first model considered was

$$Y_i = 10 - 0.9 \times log(x_i) + \epsilon_i$$

[2] "Understanding Bootstrap Confidence Interval Output." Methods. 13 Sept. 2019, blog.methodsconsultants.com/posts/understanding-bootstrap-confidence-interval-output-from-the-r-boot-package/.

[3] Wicklin, Rick. "The Bias-Corrected and Accelerated (BCa) Bootstrap Interval." *SAS Blogs*, 12 July 2017, blogs.sas.com/content/iml/2017/07/12/bootstrap-bca-interval.html.

where i was the sample size and ranged from 3 to 10, x_i represented the time in minutes and took on unique values from the exposure times in the original data (5, 10, 15, 20, 30, 40, 60, 120, 240, 360, and 480 minutes), and $\epsilon_i \sim N(0, 0.25^2)$. The β coefficients and error variance were chosen as the average of the estimated coefficients and squared residual standard error, respectively, for the eight chemicals used in this analysis. Since $\beta_1 = -0.9$, the true ten Berge exponent is $-\frac{1}{-0.9} = \frac{10}{9} = 1.11$.

For each sample size N from 3 to 10, N exposure times were randomly selected from the above list without replacement. The response values were then generated, and the regression model was fit. The ten Berge exponent was computed as $-\frac{1}{\hat{\beta}_1}$ and the five methods were run 1,000 times to generate a 95 percent confidence interval or credible interval for this value, depending on the type of method. This process was then repeated 1,000 times, and the percentage of the confidence intervals that contained the true ten Berge exponent of $\frac{10}{9}$ was computed as the coverage probability. Although the interpretation of confidence intervals and credible intervals differ, the frequentist approach of computing coverage probability is a common technique used to compare the methods. The width of each of the 1,000 intervals was also recorded, and a simple average was taken to obtain the average width.

Results

The results of this simulation are shown in Table 3 below. Note that the bootstrap resampling technique selects values from the data with replacement; with such small sample sizes, it is probable that the same value will be selected for all observations of a given run. If all values are the same, linear regression cannot be performed, so a value of 'NA' is recorded, and these values are ignored when determining the confidence intervals. The jackknife method does not have this issue, as there are always at least two unique points that can be used to fit the regression model.

Table 3: Coverage Probabilities of Intervals for Various Methods Using $n = 1.11$.

Sample Size	Jackknife (Percentile)		Bootstrap (Percentile)		Bootstrap (BCa)		Bayesian (ETI)		Bayesian (HDI)	
	CP	AW	CP	AW	CP	AW	CP	AW	CP	AW
3	0.673	4.93	0.673	4.94	0.622	4.70	1.000	1.68	1.000	1.41
4	0.588	0.51	0.919	5.18	0.895	3.80	0.989	1.09	0.990	0.95
5	0.530	0.20	0.926	0.97	0.934	1.64	0.985	0.76	0.988	0.68
6	0.514	0.13	0.926	0.52	0.950	0.65	0.968	0.60	0.969	0.56
7	0.485	0.10	0.910	0.32	0.943	0.34	0.974	0.49	0.970	0.46
8	0.444	0.08	0.896	0.27	0.929	0.27	0.978	0.43	0.977	0.41
9	0.445	0.07	0.885	0.24	0.944	0.24	0.963	0.38	0.969	0.36
10	0.429	0.07	0.900	0.22	0.942	0.22	0.968	0.34	0.969	0.33

CP: Coverage Probability
AW: Average Interval Width

At a sample size of 3, the smallest number of observations considered in this analysis to perform linear regression, the three resampling methods performed similarly, but poorly, with only 62.2 percent of the confidence intervals containing the true ten Berge exponent for the BCa bootstrap method, and 67.3 percent containing the true value for both the jackknife and the percentile bootstrap methods. Further, these intervals had an average width over 4, which is wider than the typical range. For both Bayesian approaches, all 1,000 simulated credible intervals contained the true value, and the average width was under 2. Increasing the number of observations to 4, the coverage probability for both bootstrap methods increased to around 90 percent, but the average widths were still high; the coverage probability for the jackknife approach decreased to just under 60 percent, but the average width dropped to only 0.51. The credible interval from the Bayesian method also decreased slightly, but nearly 99 percent of intervals still contained the true value, and the average widths decreased to around 1. As more observations were included, the coverage probability for the jackknife method and Bayesian method continued to decrease, on average. The coverage probabilities for the jackknife intervals dropped to under 50 percent for a sample size of 10, while the credible intervals from the Bayesian approach still contained the true exponent over 95 percent of the time. The decrease in these coverage probabilities is caused by a decrease in the width of the intervals as the sample size increases, with average widths under 0.1 for the jackknife approach and just over 0.33 for the Bayesian approaches. In contrast, the probabilities for the

bootstrap methods remained relatively stable, near 90 percent, as the sample size increased. The average widths for these intervals also decreased, becoming smaller than the Bayesian intervals at sample sizes of 6 or 7.

In all but two sample sizes considered, the HDI Bayesian method had the highest coverage probabilities, surpassed slightly by the ETI Bayesian method at sample sizes of 7 and 8. Further, the HDI Bayesian method had smaller average widths than the ETI Bayesian method for all sample sizes, as well as the bootstrap methods for small sample sizes. Thus, the HDI Bayesian approach outperformed all three resampling techniques and the ETI Bayesian method, with higher coverage probabilities and smaller intervals, indicating more reliable results.

To analyze the robustness of these intervals for different true values of β, we ran the same simulation as above using $\beta = -1.1$, which corresponds to a ten Berge exponent of 0.91, close to the lower range of expected values. Table 4 shows the coverage probabilities and average widths for the five different methods below.

Table 4: Coverage Probabilities of Intervals for Various Methods Using *n* = 0.91.

Sample Size	Jackknife (Percentile)		Bootstrap (Percentile)		Bootstrap (BCa)		Bayesian (ETI)		Bayesian (HDI)	
	CP	AW	CP	AW	CP	AW	CP	AW	CP	AW
3	0.655	7.23	0.655	7.23	0.619	7.07	0.951	1.79	1.000	1.47
4	0.610	0.29	0.934	37.53	0.922	35.57	0.945	1.15	0.986	0.96
5	0.570	0.12	0.923	0.53	0.924	4.76	0.948	0.75	0.981	0.64
6	0.508	0.09	0.933	0.29	0.959	0.34	0.933	0.54	0.965	0.48
7	0.493	0.07	0.903	0.22	0.941	0.23	0.945	0.41	0.975	0.37
8	0.478	0.06	0.892	0.19	0.937	0.19	0.952	0.34	0.978	0.31
9	0.465	0.05	0.906	0.17	0.951	0.17	0.942	0.28	0.959	0.27
10	0.426	0.05	0.895	0.15	0.932	0.15	0.952	0.25	0.969	0.24

CP: Coverage Probability
AW: Average Interval Width

From this table, we see that when the true exponent is 0.91, the coverage probabilities for the jackknife and two bootstrap methods are nearly identical to the case when the true *n* was 1.11. The average widths are slightly smaller, with the exception of the smaller sample sizes in which several

samples produced intervals with extremely wide intervals, greatly influencing the averages. For the Bayesian methods, the coverage probabilities decreased for most sample sizes, but remained over 95 percent in each case. As with the resampling techniques, the average widths of these intervals also decreased slightly when compared to a true *n* of 1.11, other than for sample sizes of 3 and 4, in which the widths increased slightly. These Bayesian approaches are based on a prior distribution where the probability that the ten Berge exponent is less than or equal to 0.91 is near 5 percent, so the credible intervals are less likely to extend much below this value; this results in slightly lower coverage probabilities, in general. Even then, for sample sizes of three through ten, the coverage probabilities are still higher for the Bayesian methods than for the other three approaches. Again, the HDI Bayesian method outperforms the ETI Bayesian method, with slightly higher coverage probabilities, and smaller or comparable average widths.

One final simulation was run, using $\beta = -0.5$, which corresponds to a ten Berge exponent of 2. This is larger than the average ten Berge exponent for the eight chemicals in this study, but still below the upper limit of expected values. The results are displayed in Table 5 below.

Table 5: Coverage Probabilities of Intervals for Various Methods Using *n* = 2.

Sample Size	Jackknife (Percentile)		Bootstrap (Percentile)		Bootstrap (BCa)		Bayesian (ETI)		Bayesian (HDI)	
	CP	AW	CP	AW	CP	AW	CP	AW	CP	AW
3	0.653	11.42	0.656	11.64	0.597	10.37	1.000	2.43	0.997	2.03
4	0.597	3.80	0.923	67.40	0.915	15.74	0.983	2.07	0.973	1.85
5	0.539	0.93	0.910	4.70	0.916	14.95	0.972	1.72	0.956	1.59
6	0.512	6.99	0.930	2.35	0.945	9.88	0.957	1.55	0.943	1.46
7	0.483	0.40	0.896	1.51	0.933	1.86	0.967	1.34	0.953	1.29
8	0.465	0.28	0.892	1.00	0.933	1.02	0.973	1.23	0.960	1.18
9	0.448	0.27	0.906	0.91	0.943	0.90	0.963	1.11	0.949	1.07
10	0.425	0.22	0.894	0.75	0.936	0.75	0.961	1.03	0.951	1.00

CP: Coverage Probability
AW: Average Interval Width

Again, each method has similar coverage probabilities as the previous two examples for all sample sizes considered; however, in this case, the ETI Bayesian method has larger coverage probabilities than all other methods, including the HDI Bayesian method. Even then, the coverage probabilities

for the two Bayesian methods differ by less than two percent, and the HDI method still has smaller average widths for all sample sizes.

Based on this simulation, a larger number of the HDI Bayesian credible intervals contain the true ten Berge exponent than the other methods when *n* is 1.11 or 0.91; at a true *n* of 2, the ETI Bayesian credible interval is slightly more accurate, but by less than two percent. For the HDI approach, over 95 percent of the intervals contained the true value for samples with three to ten observations when the true exponent was 1.11 and 0.91. When the true exponent was 2.0, closer to the upper bound of the expected range, the coverage probabilities were slightly lower, but still over 94 percent. However, in all considered cases, the HDI Bayesian method had the smallest average widths for the smallest sample sizes, and outperformed the ETI method for larger sample sizes as well.

Discussion

Balancing high coverage probabilities and low interval widths, the HDI Bayesian method is the superior method to provide insightful and accurate interval estimates for the ten Berge exponent. It outperforms the jackknife and bootstrap methods for all sample sizes in generating intervals that measure the uncertainty in the estimates for the ten Berge exponent, and performs similarly to, or slightly better than, the ETI method. This approach is most appropriate and has small, yet accurate, credible intervals when the true ten Berge exponent is closer to the center of 0.85 and 3.5, but still produces reliable results when exponents are near the edge, especially when more observations are available.

It is interesting to note that for the bootstrapping techniques, having at least four observations greatly increases the accuracy of the confidence interval. Thus, if the bootstrap method is used instead of the Bayesian method, it is recommended that at least four observations be used in order to obtain accurate interval estimates.

APPENDIX B: Model Diagnostic Plots

Figure 5: Model Diagnostic Plots for Ammonia.

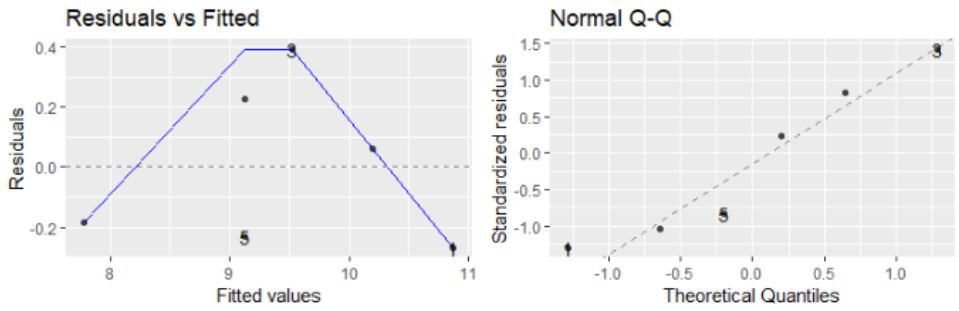

Figure 6: Model Diagnostic Plots for Acrolein.

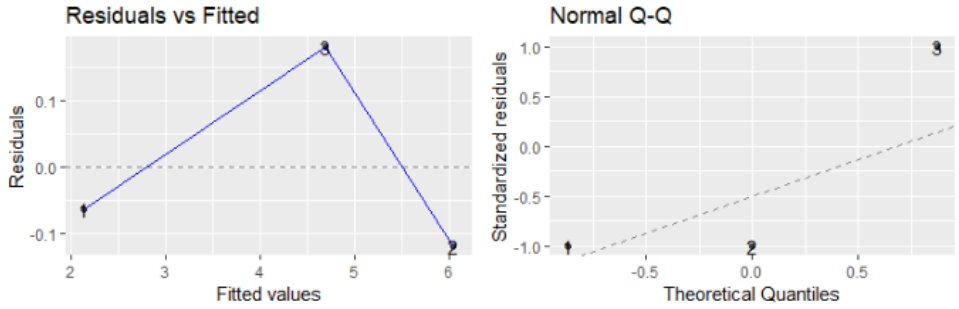

Figure 7: Model Diagnostic Plots for Carbon disulfide.

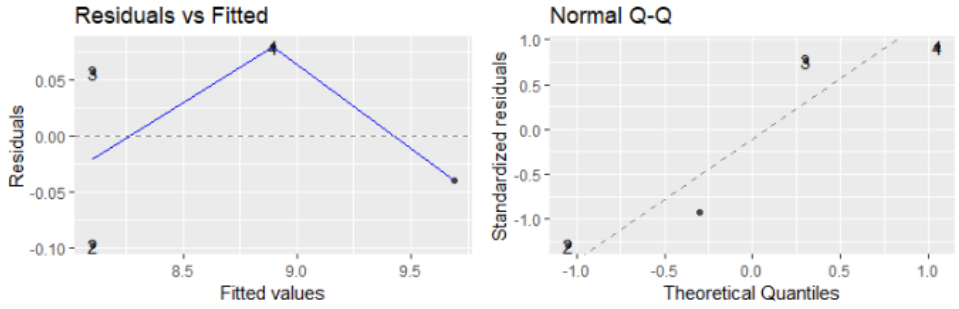

Figure 8: Model Diagnostic Plots for Carbon monoxide.

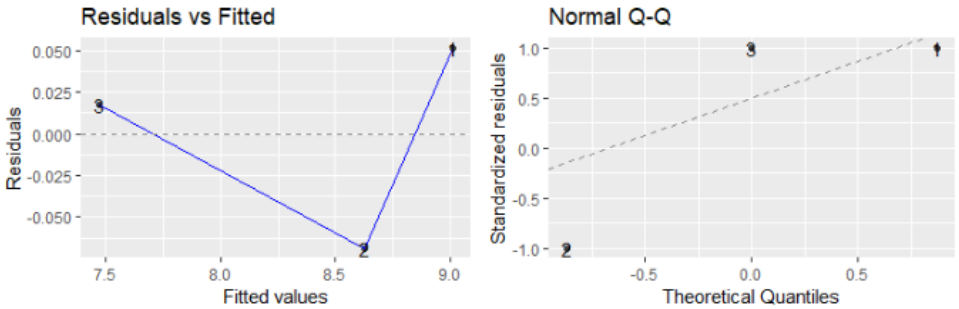

Figure 9: Model Diagnostic Plots for Epichlorohydrin.

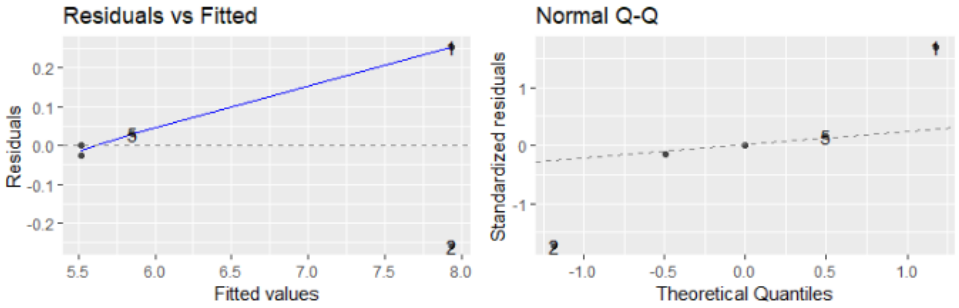

Figure 10: Model Diagnostic Plots for Ethyleneimine.

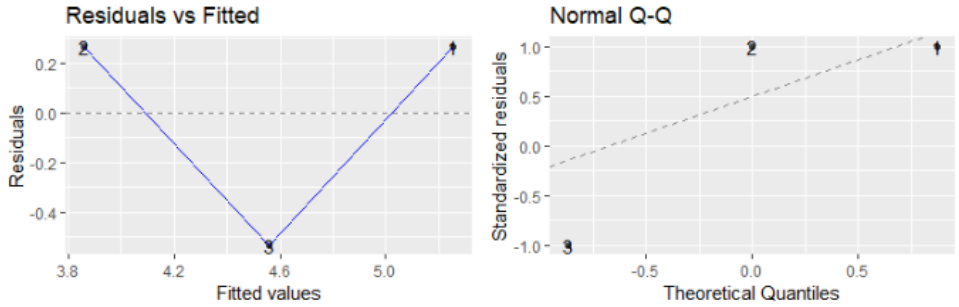

Figure 11: Model Diagnostic Plots for Hydrogen cyanide.

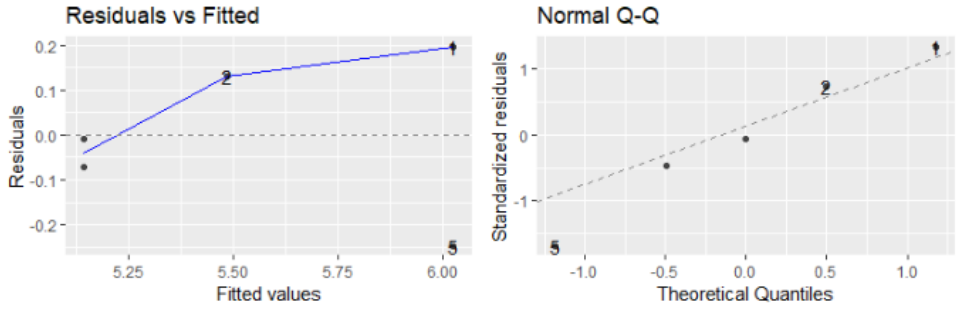

Figure 12: Model Diagnostic Plots for Pentaborane.