A SIMULATION STUDY OF THE ROBUSTNESS OF PREDICTION INTERVALS FOR AN INDEPENDENT OBSERVATION OBTAINED FROM A RANDOM SAMPLE FROM AN ASSUMED LOCATION-SCALE FAMILY OF DISTRIBUTIONS

by

NATALYA MAKAROVA

A REPORT

submitted in partial fulfillment of the requirements for the degree

MASTER OF SCIENCE

Department of Statistics
College of Arts and Sciences

KANSAS STATE UNIVERSITY
Manhattan, Kansas

2012

Approved by:

Major Professor
Dr. Paul I. Nelson

Copyright

Abstract

Suppose that based on data consisting of independent repetitions of an experiment a researcher wants to predict the outcome of the next independent outcome of the experiment. The researcher models the data as being realizations of independent, identically distributed random variables $\{X_i, i = 1, 2, ...n\}$ having density $f(\cdot)$ and the next outcome as the value of an independent random variable Y, also having density $f(\cdot)$. We assume that the density $f(\cdot)$ lies in one of three location-scale families: standard normal (symmetric); Cauchy (symmetric, heavy-tailed); extreme value (asymmetric.). The researcher does not know the values of the location and scale parameters. For $f(\cdot)$ = $f_0(\cdot)$ lying in one of these families, an exact prediction interval for Y can be constructed using equivariant estimators of the location and scale parameters to form a pivotal quantity based on $\{X_i, i = 1, 2, ...n\}$ and Y . This report investigates via a simulation study the performance of these prediction intervals in terms of coverage rate and length when the assumption that $f(\cdot)$ = $f_0(\cdot)$ is correct and when it is not.

The simulation results indicate that prediction intervals based on the assumption of normality perform quite well with normal and extreme value data and reasonably well with Cauchy data when the sample sizes are large. The heavy tailed Cauchy assumption only leads to prediction intervals that perform well with Cauchy data and is not robust when the data are normal and extreme value. Similarly, the asymmetric extreme value model leads to prediction intervals that only perform well with extreme value data. Overall, this study indicates robustness with respect to a mismatch between the assumed and actual distributions in some cases and a lack of robustness in others.

Table of Contents

List of Figures vi
List of Tables vii
Acknowledgements viii
Dedication ix
Chapter 1 - Introduction 1
  Topic of the report 1
  Literature Review and Related work 1
    Advantages and Limitations of Prediction Intervals 3
    Common applications 4
  Preliminaries 4
Chapter 2 - Constructing Prediction Intervals 6
  Prediction Interval under Normality 6
  Constructing Prediction Intervals for Arbitrary Location-Scale Families 8
  Goal of the report 9
Chapter 3 - The Simulation Study 10
  Simulation Algorithm 11
  Sensitivity Analysis of Critical Values of K Statistic 12
Chapter 4 - Results 13
  Assumed Model: Normal Distribution 13
  Assumed Model: Cauchy Distribution 14
  Assumed Model: Extreme Value Distribution 17
  Another Look at the Simulation Study 19
    Normal Distribution 19
    Cauchy distribution 20
    Extreme Values Distribution 21
Chapter 5 - Discussion 22
  Normal distribution 22
  Cauchy Distribution 22

Extreme Value Distribution 23
Another look at the simulation 23
Conclusion 24
References 25
Appendix A - Additional Tables and Figures 27
Appendix B - R Code for the Simulation Study 29

List of Figures

Figure 2.1 Simulated Independent, Standard Normal Data, n = 10. 7
Figure 4.1 Superimposed Histograms of Simulated K-statistic Based on Cauchy Data and Normal Data for small and large sample sizes (M=100000) 16
Figure 4.2 Superimposed Histograms of Simulated K-statistic Based on Cauchy Data and Extreme Value Data for small and large sample sizes (M=100000) 16
Figure 4.3 Relative Frequency Histogram of Sample Standard Deviation from Cauchy Data for small and large sample sizes (M=100000).. 17
Figure 4.4 Superimposed Histograms of Simulated K-statistic Based on Extreme Values Data and Normal Data for small and large sample sizes (M=100000) (M=100000).................... 19
Figure A.1 Estimated Coverage Rates and 5% Confidence Intervals. Cauchy Assumptions 28

List of Tables

Table 4.1 Estimated Average Lengths and Coverage Rates Assuming Normality, L = 1000..... 13

Table 4.2 Estimated Lengths and Coverage Rates Assuming Cauchy Distribution, M=100000, L = 1000 14

Table 4.3 Estimated Average Lengths and Coverage Rates Assuming Extreme Value Distribution, M=100000, L = 1000 18

Table 4.4 Estimated Average Lengths and Coverage Rates for Normally Distributed Data 20

Table 4.5 Estimated Average Lengths and Coverage Rates for Cauchy Distributed Data.......... 20

Table 4.6 Estimated Average Lengths and Coverage Rates for Extreme Value Distributed Data 21

Table A.1 Results Regarding Sensitivity of Critical Values of K statistics from Cauchy Distribution 27

Acknowledgements

It is a pleasure to express my profound thanks I feel for all of those people who helped me along the way. First and foremost, I express my immense gratitude to my major advisor Dr. Paul I. Nelson for his support and encouragement, for his patience during my work on this report. I would also like to thank Dr. Abigail Jager and Dr. Gary Gadbury for their time and service on my committee.

Great thanks to my academic teachers here at Kansas State University who I had privilege to have as my instructors: Dr. Nora Bello, Dr. Gary Gadbury, Dr. Weixin Yao, Dr. Abigail Jager , Dr. Leigh Murray, Dr. Haiyan Wang, Dr. Weixing Song, and Dr. Kun Chen. Thank you for being patient with the numbers of questions I asked. You are all are cheerful and dedicated teachers and mentors.

Appreciation also goes to all graduate students I study with for making my graduate school a meaningful and exciting experience. My special thanks go to Chris Juarez, Garth Highland, Indu Seetharaman and Beth Larrabee for willingness to share their ideas and for providing their valuable help throughout two years at Kansas State University.

I would like to thank my family, my husband Aleksandr Elesev, and my parents Ludmila Makarova and Aleksandr Makarov, for supporting and encouraging me to pursue this degree.

Dedication

I am dedicating this report to my loving husband, Aleksandr Elesev, who stood by my side and encouraged me to pursue this degree. Thank you for being patient with the lack of time available during the last two years. Thank you for being with me when I was in doubt. This is as much your degree as it is mine.

Chapter 1 - Introduction

This chapter develops the motivation for studying the robustness of prediction intervals, provides a literature review, compares prediction intervals to confidence intervals and introduces some important terms and definitions.

Suppose that responses $\underline{\boldsymbol{x}} = \{x_i; i = 1,2, \dots, n\}$, the available data, are obtained from independent repetitions of an experiment-process carried out under similar conditions and it is desired to *predict* the response Y that would be obtained if the experiment-process were independently carried out again. This setting does not include the use of covariates and hence does not fall under the usual heading of prediction in regression models. A prediction may be made by constructing an interval estimate $PI(\underline{X})=(L(\underline{X}),U(\underline{X}))$ and being able to assign some form of likelihood to the statement that Y lies in $PI(\underline{X})$. For example, suppose that $\underline{\boldsymbol{x}}$ consists of the yields per acre of variety V wheat planted on n fields by a farmer. Rather than estimating the mean yield of variety V wheat per acre across all such fields, the farmer might very well be interested in using $PI(\underline{X})$ to predict the actual yield per acre she will obtain the next time she plants variety V. For another example, consider a person just diagnosed with lung cancer. A confidence interval for the mean survival time of all such newly diagnosed patients only incorporates the uncertainty arising from using a point estimator of the mean and fails to fully take into account the variability of lifetimes around their population mean. A prediction interval is based on estimates of both of these sources of variability and is accordingly typically wider than the corresponding confidence interval, which understates the uncertainty of current knowledge about the patient's future survival time.

Topic of the report

An assessment of the performance in terms of width and coverage rate of parametric prediction intervals for a future independent observation from an assumed location-scale family of distributions when that assumption is true and when it is false

Literature Review and Related work

Aitchison and Dunsmore (1975) give examples of the use of prediction intervals in a variety of settings. Christoffersen (1998) gives some examples from economics and finance.

Although confidence intervals for a mean are often relatively easy to compute and interpret, Christoffersen argues that interval prediction is a better tool than interval parameter estimation for economic planning. For example, a central bank governor would be more interested in forecasting the actual inflation rate over the next six months than estimating its mean inflation in order to carry out a monetary policy. A production manager planning to purchase inventory needs to predict sales in order to decide how much to order. A prediction interval might serve as a control bound for assessing product quality on an assembly line. Data points that are beyond the control bounds are referred to as being “out of control” and could indicate that remedial action needs to be taken. Patel (1989) also states that prediction intervals could provide guidelines for establishing warranty limits for the future performance of a product.

A large amount of research has been done on prediction. In particular, the concept of a predictive distribution has been discussed over a long period of time, starting with Laplace’s (1814) attempt almost two hundred years ago to calculate the probability of obtaining a success in a future Bernoulli trial based on prior information. Patel (1989) states that one of the early papers on prediction intervals is Baker (1935), which derived the probability density function of a deviation from the mean that would occur in a future sample based on the information from the observed sample. Since a prediction interval is a special case of a tolerance interval, additional references are to be found in the literature on tolerance intervals. In addition to Aitchison and Dunsmore (1975) a number of works in the latter part of the twentieth century presented results on the exact and approximate derivation of prediction intervals for a variety of distributions and settings. Patel (1989) gives a comprehensive review of those results. He also provides a long list of related research. Chatfield (1993) emphasizes the importance of prediction intervals, their advantages and limitations, and provides a summary of different methods of constructing prediction intervals for time series. Geisser (1993) is an important monograph on predictive distributions and the Bayesian approach to prediction. Despite the large amount of publications on *PI*s and their practical importance, *PI*s usually appear in textbooks only in the context of regression. Abraham and Ledolter (1983) is a notable exception. My initial literature search didn’t reveal any work regarding the robustness of prediction intervals for a future independent observation from an assumed location-scale family of distributions, the topic of my report. Olive (2007), (2003) and Fisher and Horn (1994) do consider the problem of constructing prediction intervals in a regression setting when normality does not hold.

Advantages and Limitations of Prediction Intervals

There are a number of advantages of interval prediction over a point forecast and a confidence interval for the mean:

- Prediction intervals are especially useful because they can predict what a future value, such as the height of a river subject to flooding, is *likely* to be before it happens. A confidence interval for the mean in this case would only estimate mean height across a long time period.
- A main advantage of a prediction interval over a point estimate is that it takes into account the variation of the future observation around the point estimate. The majority of individual future outcomes deviate from the point estimate, which a *PI* takes into account. Prediction intervals allow different strategies to be used to accommodate a variety of possible outcomes.
- Unlike parameter estimation, for instance confidence interval for mean (*CI*), which makes statements about hypothetical quantities, such as population means, that can almost never be verified, predictions can in principle be checked by observing what happens.

There are some limitations of using *PI* which one should consider:

- The prediction interval might be rather wide when the process under study produces outcomes that have large variation. In some situations the large width of a prediction interval limits its usefulness in decision making. If, for example, the weather forecaster predicts with 95% ‘confidence’ that the temperature tomorrow will range from 15F to 100F ,we can’t make a decision whether to wear a coat or a t-shirt tomorrow.
- With increasing sample size, the widths of prediction intervals decrease but do not converge to zero, as happens with confidence intervals for the mean.
- Theoretical *PI*s are difficult or impossible to construct, particularly for data coming from complex distribution or for complex nonlinear multivariate models. In this case simulation procedures have been used to construct approximate PI’s.

Common applications

- Regression. Assuming normality with constant variance, the least squares fitted model and variance estimate can be used to construct exact prediction intervals for specified levels of the independent variable(s).
- Time series is another traditional area for prediction. These models are widely used by forecasters in Economics and Finance.
- Nonparametric prediction intervals. This approach is helpful when no assumptions about the initial distribution can be made. Here, simulation helps to construct the *PI*.
- Bayesian prediction intervals.

Preliminaries

Prediction Interval

Let $\underline{\boldsymbol{x}} = \{x_i; i = 1,2,\dots,n\}$ be the set of observed "past" values of random variables $\underline{X} = \{X_i; i = 1,2,\dots,n\}$ in an experiment carried out in a series of independent trials and *Y* be an unobserved "future" value. The random variables $\{X_i\}$ and *Y* are by design independent, all having the same distribution, *F(x)*. We say that interval $PI(\underline{X})=(L(\underline{X}),U(\underline{X}))$ is a (1-α) two-sided prediction interval for *Y* if $Pr\left[Y \in PI(\underline{X})\right] = 1 - \alpha$. In other words, a $(1-\alpha)$ prediction interval for *Y* is an interval determined from past observations such that the probability that the "future" *Y* will fall in the interval is equal to $1-\alpha$. Suppose exact $(1-\alpha)$ prediction intervals for a "future" *Y* are independently constructed for many such pairs of ($\underline{X}$) and *Y*. Then, about 100(1-α)% of these intervals will contain the corresponding *Y* .

As noted above, prediction intervals are generally wider in length than corresponding confidence intervals because, unlike confidence intervals, prediction intervals account for two sources of variability. First, unknown parameters need to be estimated using the observed data. This is the only source of variation that a confidence interval for a population mean needs to incorporate. Second, in prediction, variation of the future observation about its mean needs to be accounted for. The performance of the prediction intervals studied in this report will be assessed in terms of length and coverage rate.

Length of Prediction Interval

Length of the prediction interval $PI(\underline{X})=(L(\underline{X}),U(\underline{X}))$ is $(U(\underline{X})- L(\underline{X}))$.

Coverage rate of Prediction Interval

Under assumed distributional forms, this report will construct prediction intervals $PI(\underline{X})=(L(\underline{X}),U(\underline{X}))$ such that $Pr[Y \in PI(\underline{X})] = 1-\alpha$ when these assumptions are valid and use simulation to investigate $Pr[Y \in PI(\underline{X})]$, called the (actual) *coverage rate* or coverage probability, when the assumptions do not hold. . The target coverage rate $1-\alpha$ will also be referred to as the *nominal coverage rate*. Large differences between actual coverage rates and nominal coverage rates constitute poor performance and what we term *a lack of robustness*. We are going to compare the nominal (assumed) coverage rate and observed (coming from multiple simulations) coverage rates.

Chapter 2 - Constructing Prediction Intervals

This chapter develops a procedure for constructing prediction intervals from known location-scale families of distributions, presents a formal description of the report's goal and explains the initial settings for computation.

Prediction Interval under Normality

Starting with the familiar assumption of normality, suppose we have observed a random sample $\underline{\boldsymbol{x}} = \{x_i; i = 1,2, \ldots, n\}$ and want to predict a "future", unobserved random value Y, where all the random variables are iid from $N(\mu, \sigma^2)$, a normal distribution with mean μ and variance σ^2, both unknown. Let $\bar{X} = \frac{1}{n}\sum_{i=1}^{n} X_i$ and $S = \sqrt{\frac{\sum_{i=1}^{n}(X_i - \bar{X})^2}{n-1}}$ be the sample mean and sample standard deviation. Using standard results, we have that $Z = \frac{\mathrm{Y}-\bar{\mathrm{X}}}{\sigma\left(1+\frac{1}{\mathrm{n}}\right)^{\frac{1}{2}}} \sim \mathrm{N}(0,1)$, independent of $\sqrt{\frac{\sum\left(X_i - \bar{X}\right)^2}{\mathrm{n}-1}}$, and $\omega = \frac{(\mathrm{n}-1)\mathrm{S}^2}{\sigma^2} \sim \chi^2$ with n-1 degrees of freedom. Since $\frac{Z}{\sqrt{\frac{\omega}{n-1}}} \sim \mathrm{t}_{\mathrm{n}-1}$ by the definition of *t*-distribution with (*n-1*) degrees of freedom, this yields

$K = \frac{\mathrm{Y}-\bar{\mathrm{X}}}{\mathrm{S}} = \frac{\sqrt{\frac{n+1}{n}\sigma^2}Z}{\sqrt{\frac{\omega}{n-1}\sigma^2}} = \sqrt{\frac{n+1}{n}}\frac{Z}{\sqrt{\frac{\omega}{n-1}}} \sim \mathrm{t}_{\mathrm{n}-1}\sqrt{\frac{n+1}{n}}$, where $\mathrm{t}_{\mathrm{n}-1}$ denotes a *t*-distribution with $\mathrm{n}-1$ degrees of freedom. Then, an the <u>exact</u> (*1-α*) prediction interval for *Y* is obtained by re-expressing a probability statement about t_{n-1} in terms of a statement about *Y*, as follows:

$$1 - \alpha = \Pr\left(t_{\frac{\alpha}{2},df=n-1} \leq \frac{Z}{\sqrt{\frac{\omega}{n-1}}} \leq -t_{\frac{\alpha}{2},df=n-1}\right) =$$

$$= \Pr\left(\sqrt{\frac{n+1}{n}} * t_{\frac{\alpha}{2},df=n-1} \leq \sqrt{\frac{n+1}{n}} * \frac{Z}{\sqrt{\frac{\omega}{n-1}}} \leq -\sqrt{\frac{n+1}{n}} * t_{\frac{\alpha}{2},df=n-1}\right) =$$

$$= \Pr\left(\sqrt{\frac{n+1}{n}} * t_{\frac{\alpha}{2},df=n-1} \leq \frac{\mathrm{Y}-\bar{\mathrm{X}}}{\mathrm{S}} \leq -\sqrt{\frac{n+1}{n}} * t_{\frac{\alpha}{2},df=n-1}\right) =$$

$$= \Pr\left(\bar{X} - S * \sqrt{\frac{n+1}{n}} * t_{1-\frac{\alpha}{2},\ df=n-1} \leq Y \leq \bar{X} + S * \sqrt{\frac{n+1}{n}} * t_{1-\frac{\alpha}{2},\ df=n-1}\right)$$

Specifically, the prediction interval is given by

$$\bar{x} \mp t_{\frac{\alpha}{2},n-1} S\left(\sqrt{\frac{n+1}{n}}\right) \qquad (1)$$

Note that this prediction interval differs from the corresponding confidence interval for μ by the *stretching* factor $\sqrt{((n+1)/n)}$.

Example

Simulated data of n = 10 observations from the normal theory model $N(\underline{0}, 1)$ are marked in blue in Figure 2.1 below. The value to be predicted, 'y' is marked in 'red.'

Figure 2.1 Simulated Independent, Standard Normal Data, n = 10.

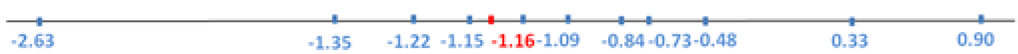

Evaluating the prediction interval given in the (1), these data leads to the 0.95 prediction interval

$$-0.83 \mp t_{0.025,10-1} * 0.96 * \sqrt{1+\frac{1}{10}} \quad \text{or} \quad [-3.11,\ 1.46]$$

Note that "future" observation Y = -1.16 does fall inside the interval and that a .95 confidence interval for μ, which is zero here, is given by $\bar{x} \mp t_{\frac{\alpha}{2},n-1}\frac{S}{\sqrt{n}} = -0.83 \mp 2.262 * \frac{0.96}{\sqrt{10}}$ or *[-1.52, -0.14]*. Notice here that zero, the mean, does not fall into the confidence interval, a coverage failure which occurs with probability 5% . As noted above, the prediction interval for a "future" observation is larger than confidence interval for the population mean.

Since the distribution of K is free of μ and σ^2, this approach can be extended to models where the data are generated from any assumed location-scale family. However, outside of normality, the exact distribution of K would be difficult to obtain in a useful, closed form. Therefore, instead of even attempting exact derivations, in the next section we will demonstrate how to use simulation to approximate quantiles of the distribution of K and use them to construct prediction interval for several distributions.

Constructing Prediction Intervals for Arbitrary Location-Scale Families

Responses $\underline{\boldsymbol{x}} = \{x_i; i = 1,2, \dots, n\}$ are assumed to be realized values of independent, identically distributed random variables $\underline{X} = \{X_i; i = 1,2, \dots, n\}$, each having absolutely continuous, partially specified distribution function *F(x)* and *Y* is independent of $\underline{X}$ and also has distribution function $F(x)$. Further, assume that for unknown location parameter μ and unknown scale parameter $\sigma > 0$,

$$F(x) = F_0((x-\mu)/\sigma) \text{ ,}$$

where F_0 is fully specified. Let, $\hat{\mu}(\underline{X})$ be an estimator of μ and $\hat{\sigma}(\underline{X})$ an estimator of σ that are equivariant in the sense that for any constants $a > 0$ and b

$$\hat{\mu}\left(a\underline{X} + b\underline{1}\right) = a\hat{\mu}\left(\underline{X}\right) + b \text{ ,}$$

$$\hat{\sigma}\left(a\underline{X} + b1\right) = a\hat{\sigma}(\underline{X}).$$

For example, $\hat{\mu}(\underline{X})$ could be the sample mean or median and $\hat{\sigma}(\underline{X})$ could be the sample standard deviation or sample inter-quartile range. Let $Z_x = (X - \mu)/\sigma$, $Z_y = (Y - \mu)/\sigma$, and note that $Pr(Z_x < z) = Pr\left(Z_y < z\right) = F_o(z)$. Then, the statistic

$$\begin{aligned} K(F_0, n) &\equiv (Y - \hat{\mu}(\underline{X}))/\hat{\sigma}(\underline{X}) \\ &= (Z_y - \hat{\mu}(Z_x))/\hat{\sigma}(Z_x) \end{aligned} \qquad (2)$$

has a distribution determined by F_0 and hence, at least in theory, which can be computed.

To use equivariance to construct a prediction interval $PI(\underline{X})$ with $Pr\left[Y \in PI(\underline{X})\right] = 1 - \alpha$, find quantiles $k_{\frac{\alpha}{2}}(F_0 \,, n)$ and $k_{1-\frac{\alpha}{2}}(F_0 \,, n)$, henceforth called critical values, where for any $\delta \in (0,1)$, $\Pr\left(K(F_0, n) \le k_\delta(F_0, n)\right) = \delta$. Then,

$$\Pr\left(k_{\frac{\alpha}{2}}(F_0 \,, n) \le K(F_0, n) \le k_{1-\frac{\alpha}{2}}(F_0 \,, n)\right) = 1 - \alpha.$$

Hence, having observed $\underline{\boldsymbol{x}}$, an exact $1 - \alpha$ prediction interval for *Y* is then given by

$$\hat{\mu}\left(\underline{\boldsymbol{x}}\right) + k_{\frac{\alpha}{2}}(F_0 \,, n)\hat{\sigma}\left(\underline{\boldsymbol{x}}\right),\ \hat{\mu}\left(\underline{\boldsymbol{x}}\right) + k_{1-\frac{\alpha}{2}}(F_0 \,, n)\hat{\sigma}\left(\underline{\boldsymbol{x}}\right)] \text{ .} \qquad (3)$$

When F_0 is standard normal, $\hat{\mu}(\underline{\boldsymbol{x}})$ is the sample mean and $\hat{\sigma}(\underline{\boldsymbol{x}})$ is the sample standard deviation, the critical values can be obtained, as illustrated above from the *t*-tables with degrees of freedom (*n-1*). In other cases they can be approximated by simulation, as described below.

Goal of the report

This report uses a simulation study to investigate the following question: In terms of coverage rate and length, how does the interval in (3) perform when $Pr(Z_x < z) = Pr(Z_y < z) = G(z)$, where $G = F_0$ and $G \neq F_0$ The parameter settings and the simulation design are described below.

Three representative choices for $f_0 \equiv F_0'$ for the assumed density were used:

(1) Standard Normal: $f_0^{(1)}(x) = e^{-x^2/2}/\sqrt{2\pi}$.
(2) Cauchy: $f_0^{(2)}(x) = (1/c\pi)/[1 + (x/c)^2]$, c = 1.35/2.
(3) Extreme value: $f_0^{(3)}(x) = (1/c)e^{x/c}\exp(-e^{x/c})$,

where c = 1.35/(log(log(.25)/log(.75))). (Here Log = Ln)

In this report, the sample location estimator $\hat{\mu}(\underline{X})$ was taken to be the sample mean and the sample scale estimator $\hat{\sigma}(\underline{X})$ was taken to be the sample standard deviation.

Notes:

(i) In cases (2) and (3), the constant c is defined so that these distributions have the same inter-quartile range as a standard normal distribution. The *shape* parameter for the extreme value distribution is set at zero. This extreme value distribution is also called the Gumbel distribution.

(ii) Without loss of generality, we set $\mu = 0$ and $\sigma = 1$ in all cases.

(iii) There are many routines for generating approximate normal random variables. To generate Z from (2) or (3), I will use R to generate W from the 'standard' version, corresponding to c = 1 and set Z = cW.

(iv) I used representative values of sample size *n* ; small, medium and large.

(v) The nominal coverage rate was set $1\text{-}\alpha = 0.95$ in all cases.

(vi) I used some preliminary simulations , as described below, to specify *M* , the number of samples used to approximate the critical values used in constructing the prediction intervals and *L* = number of *PI*'s constructed under a given set of parameter values.

Chapter 3 - The Simulation Study

This chapter contains a description of the algorithms and software that were used to carry out the simulation study. I use the software package *R (www.r-project.org)*, version 2.13.00 (2011-04-13), to perform the necessary computations. *R* is a free software package that can be used for statistical computing and producing graphics displays. It compiles and runs on a wide variety of UNIX , Windows and MacOS platforms. The *R* code I used appears in Appendix B. For simulation I use the **evd** package and, in particular, built-in functions:

- **rnorm ()** – produces **n** normally distributed random variables with specified mean standard deviation. In particular I used parameters **mean**=0, **sd=**1;
- **rcauchy()** – produces **n** random variables coming from Cauchy distribution with specified parameters. In particular I used parameters **location** = 0, **scale** = 1.35/2;
- **rgev()** – produces **n** random variables coming from extreme value distribution with specified parameters. In particular, I used parameters **loc**=0, **scale=-**1.35/log(log(0.25)/log(0.75)), **shape=**0.

My simulation was carried out as follows. When F_0 is the standard normal distribution, *K* has a scaled *t*-distribution and the critical values $k_{\frac{\alpha}{2}}(F_0\,, n)$ and $k_{1-\frac{\alpha}{2}}(F_0\,, n)$ can be obtained from the *t*-table. When F_0 is the Cauchy or extreme value distribution, deriving the critical values is not feasible. Instead, I used simulation to approximate them. For each of *M* simulated samples from the standard Cauchy and extreme value distributions, the statistic *K* was calculated. Next, the estimated critical values $k_{\frac{\alpha}{2}}(F_0\,, n)$ and $k_{1-\frac{\alpha}{2}}(F_0\,, n)$ are defined as the $\alpha/2$ and $1- \alpha/2$ sample quantiles of the simulated *K*’s. Formula (3) was then used to construct the prediction intervals for data obtained as *L* independent simulations for each parameter setting. In all cases, I tallied the fraction out of the *L* data sets that contained the simulated, “future” value *Y*, a mean and median interval width. Specifically, I used the following algorithm to carry out my simulation study.

Simulation Algorithm

(a) Set index ‘ j’ =1.

(b) Find or estimate the critical values $k_{\frac{\alpha}{2}}(F_0\,,n)$ and $k_{1-\frac{\alpha}{2}}(F_0\,,n)$. For j = 1, the critical values can be obtained from the *t*-tables with degree of freedom n-1. For j = 2,3, generate iid random variables $\{Z_{il}; i = 1,2,\dots,n+1; l = 1,2,..,M\}$ from $f_0^{(j)}$, M a large number. Let K_l be $K(F_0, n)$ in (2) computed from $\{Z_{il}: i = 1,2,\dots,n+1\}$. Use sample quantiles of $\{K_l; l = 1,2,\dots,M\}$ to estimate $k_{\frac{\alpha}{2}}(F_0^{(j)}\,,n)$ and $k_{1-\frac{\alpha}{2}}(F_0^{(j)}\,,n)$.

(c) Generate the data and “future” value to be predicted $\{Z_i^{(m)}; i = 1,2,\dots,n+1\}$ from each of the three distributions $f_0^{(m)}$ and make $\{X_i^{(m)} = Z_i^{(m)}; i = 1,2,\dots,n\}$. Use $k_{\frac{\alpha}{2}}(F_0\,,n)$, and $k_{1-\frac{\alpha}{2}}(F_0\,,n)$ estimated in (b) and (3) to construct a nominal $1-\alpha$ prediction interval for $Y_i^{(m)} = Z_{n+1}^{(m)}$, m = 1,2,3 . Note that of the three prediction intervals created, only the one with m =j is ‘correct.’

(d) Check if the generated “future” observations falls inside the *PI* and record its length for each of three data sets

(e) Independently repeat (c) and (d) $L =$ 1000 times.

(f) Compute estimated coverage rates, mean widths and median widths for each of the m = 1,2,3 types of data.

(g) Go to (h) if $j = 3$. Otherwise, replace j by j +1 and repeat (b)-(f) .

(h) Independently carry out (a) – (g) for all selected sample sizes.

Before carrying out my simulation study, I constructed a preliminary simulation, described below, to investigate the sensitivity of the output of widths and coverage rates to specification of M, the number of data sets used to estimate critical values for the intervals constructed when assuming the Cauchy and extreme value models. Specifically I had to estimate the tail quantiles $k_{\frac{\alpha}{2}}(F_0\,,n)$ and $k_{1-\frac{\alpha}{2}}(F_0\,,n)$ of the distribution of K when F_0 is a standard Cauchy or extreme value distribution

Sensitivity Analysis of Critical Values of K Statistic

I carried out the algorithm given above for various choices of *M*. For each *M* used, I generated *M* independent copies of *K* and used their sample quantiles to estimate the required critical values. Table A.1 in Appendix A presents the estimated 2.5th and 97.5th percentiles of K for several choices of *M* when F_0 is from the Cauchy distribution described in the Chapter 2. The critical values are obtained for $1 - \alpha$ =0.95 and sample sizes *n*=10 and *n*=200. One may notice, that for the small sample size, n =*10***,** the critical values are quite sensitive to the number of iterations *M* used, until *M*=21000. After this point, the critical values do not change much when the number of iterations varies from 21000 to 500000. A similar pattern holds for the largest sample size used in this report, *n*=200**.** As an additional check on sensitivity, recalling that the exact critical values are symmetric about zero for the Cauchy distribution, note that in Table A.1 the estimated critical values are reasonably close to being symmetric about zero for *M* at least 7000. Based on this analysis, balancing computing time (displayed in the table A.1 as well) and accuracy, I decided to use *M* = 100000. I also used *M* = 100000 for estimating critical values based on the extreme value model.

Chapter 4 - Results

Using the simulation results, I summarize, assess and compare how well the prediction intervals perform when the assumed F_0 is correct and when it is not. Recall, that I used three distributions: normal, Cauchy and extreme value. I recorded the length of each of the L intervals for each parameter setting and summarized these lengths using their means and medians. Letting $\{Y_i\}$ denote the simulated future observations generated at each parameter setting, the estimated coverage rate for each parameter setting is given by

$$\hat{p} = \sum_{i=1}^{L} I[Y_i \in PI_i]/L$$

which has a standard error no larger than $\sqrt{0.5(0.5)/L} = 0.5/\sqrt{L}$. Then, an approximate, large sample 0.95 confidence interval for the actual coverage rate p is given by $\hat{p} \mp 1.96\sqrt{\hat{p}(1-\hat{p})/L}$. For each of the possible choices of F_0, I now present my results in the form of tables giving estimated coverage rates and 'average' interval widths. Coverage rates which appear to be far from 0.95 in a practical sense are highlighted.

Assumed Model: Normal Distribution

Table 4.1 summarizes estimated coverage rates and estimated mean and median widths when the prediction intervals are constructed assuming normality. Using the variance of a binomial distribution, the standard errors of the coverage rates in this section are no larger than 0.016

Table 4.1 Estimated Average Lengths and Coverage Rates Assuming Normality, L = 1000

1-α	n	Critical Values (scaled t)	Average Width / Median Width / Coverage Rate		
			Normal	Cauchy	Extreme Value
0.95	5	{-3.05, 3.05}	5.71/ 5.56 /0.9460	31.47 / 10.46 / 0.8960	6.28 / 5.72 / 0.9390
	10	{-2.37, 2.37}	4.53 / 4.48 / 0.9460	56.89 / 12.18 / 0.9050	4.90 / 4.73 / 0.9400
	50	{-2.03, 2.03}	4.05 / 4.05 / 0.9450	157.97 / 23.68 / 0.9550	4.39 / 4.37 / 0.9500
	200	{-1.98, 1.98}	3.95 / 3.94 / 0.9610	381.38 / 46.77 / 0.9660	4.34 / 4.32 / 0.9570

From Table 4.1 we see that the attained coverage rate is quite close to the nominal 0.95 value, except for Cauchy data with small samples sizes, n = 5 and 10. Before interpreting interval widths here and in what follows, keep in mind that the interquartile ranges (IQR's) of all of our models and data are 1.35, the same as the IQR of the standard normal, which provides a basis for determining large and small lengths. In Table 4.1, mean widths for normal and extreme value data are similar, being slightly smaller for normal data, both close to the difference between the normal model critical values, which approach ± 1.96 as sample size increases. This behavior was expected for normal data but surprising for extreme value data. For normal and extreme data, median widths are close to mean widths. For Cauchy data, median widths are more stable and much smaller than mean widths, which actually increase as sample size increases, due, no doubt, to the very heavy tails of the Cauchy distribution.

Assumed Model: Cauchy Distribution

Table 4.2 below summarizes coverage rates and lengths when the Cauchy model is assumed and M = 100000 iterations are used to estimate the critical values.

Table 4.2 Estimated Lengths and Coverage Rates Assuming Cauchy Distribution, M=100000, L = 1000

$1-\alpha$	n	Critical Values	Average Width / Median Width / Coverage Rate		
			Normal	Cauchy	Extreme Value
0.95	5	{-6.65, 6.71}	12.46 / 12.08 / 0.9960	82.01 / 22.18 / 0.9520	13.47 / 12.58 / 0.9910
	10	{-4.21, 4.38}	8.37 / 8.28 / 0.9980	202.10 / 21.96 / 0.9480	8.87 / 8.54 / 0.9900
	50	{-1.79, 1.83}	3.58 / 3.57 / 0.9080	69.75 / 20.35 / 0.9480	3.93 / 3.88 / 0.9190
	200	{-0.86, 0.91 }	1.76 / 1.76 / 0.6160	175.82 / 20.73 / 0.9400	1.94 / 1.94 / 0.6450

As expected, the actual coverage rates in Table 4.2 are close to the nominal 0.95 for Cauchy data since the Cauchy critical values are correct here. However, even with coverage rates close to their 0.95 nominal value, the lengths of the prediction intervals are large. This results, as clearly shown in Figure 4.3, results from the very right-skewed distribution of sample standard deviations obtained from Cauchy data. Specifically, for example, for n = 200, but not for n = 10, a high proportion of sample standard deviations exceed twenty (30.6% of sample standard deviations for n=200 comparing to 7% for n=10), a value much larger than the interquartile range of the distribution. Note that median lengths for the intervals constructed from Cauchy data are more stable and much smaller than mean lengths.

For normal and extreme value data, the over coverage for n=5 and n=10 and under coverage for n = 50 and 200 seen in Table 4.1 are caused by the Cauchy critical values being too far apart for small samples and too close for large samples, behavior explained in Figure 4.1 below, resulting in prediction intervals that are respectively too wide and too narrow. This over/under coverage is even clearer from Figure A.1. in the Appendix, which presents 95% confidence intervals for actual coverage rate, constructed using the binomial distribution so that confidence interval widths depend on estimated coverage rates. None of these confidence intervals for normal data includes the nominal 0.95 coverage rate. Specifically, from Table 4.2, in the case of data sampled from the normal distribution, the estimated coverage rate is 0.9080 for n=50 and only 0.6160 for n=200. Similarly, when data are sampled from the extreme value distribution, prediction intervals constructed using the Cauchy distribution had an estimated coverage rate of 0.9190 for n=50 and only 0.6450 for n=200. Under/over coverage for extreme value data using the Cauchy model critical values is also illustrated in Figure 4.2 below and Figure A.1 in Appendix.

Figure 4.1 and 4.3 graphically illustrates the reason, noted above, for these low/high coverage rates for prediction intervals constructed from Cauchy model critical values and applied to normal and extreme value data: critical values computed from the Cauchy model are too small/too large in absolute value when used with normal and extreme value data. Specifically, Figure 4.1 displays superimposed histograms of simulated values of the *K*-statistic based on data from the Cauchy distribution (green) and data from the normal distribution (yellow), where K has a scaled t-distribution. Histograms are displayed for small and large sample size. The 2.5th and 97.5th sample quantiles are marked for both cases. From Figure 4.1 we see that for the large sample size, critical values based on the Cauchy model are much closer together than those based on the normal model, resulting in prediction intervals that are too narrow for normal data and consequently have lower than nominal coverage rates. For the small sample size n=10 critical values based on the Cauchy model are much further away from one another than those based on the normal model, resulting in prediction intervals that are too wide for normal data and consequently have much higher than nominal coverage rates. Some of the extreme *K*'s obtained from Cauchy data were deleted in order to create a more detailed plot.

Figure 4.1 Superimposed Histograms of Simulated K-statistic Based on Cauchy Data and Normal Data for Small and Large Sample Sizes (M=100000)

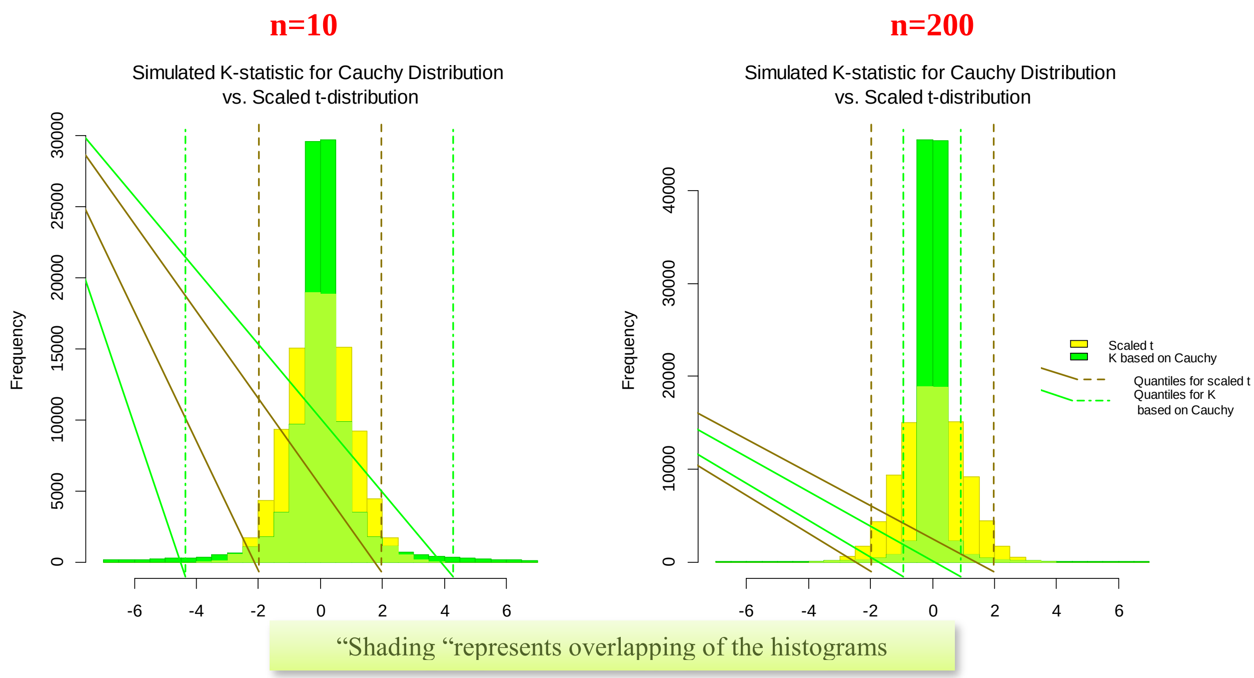

Figure 4.2 Superimposed Histograms of Simulated K-statistic Based on Cauchy Data and Extreme Value Data for Small and Large Sample Sizes (M=100000)

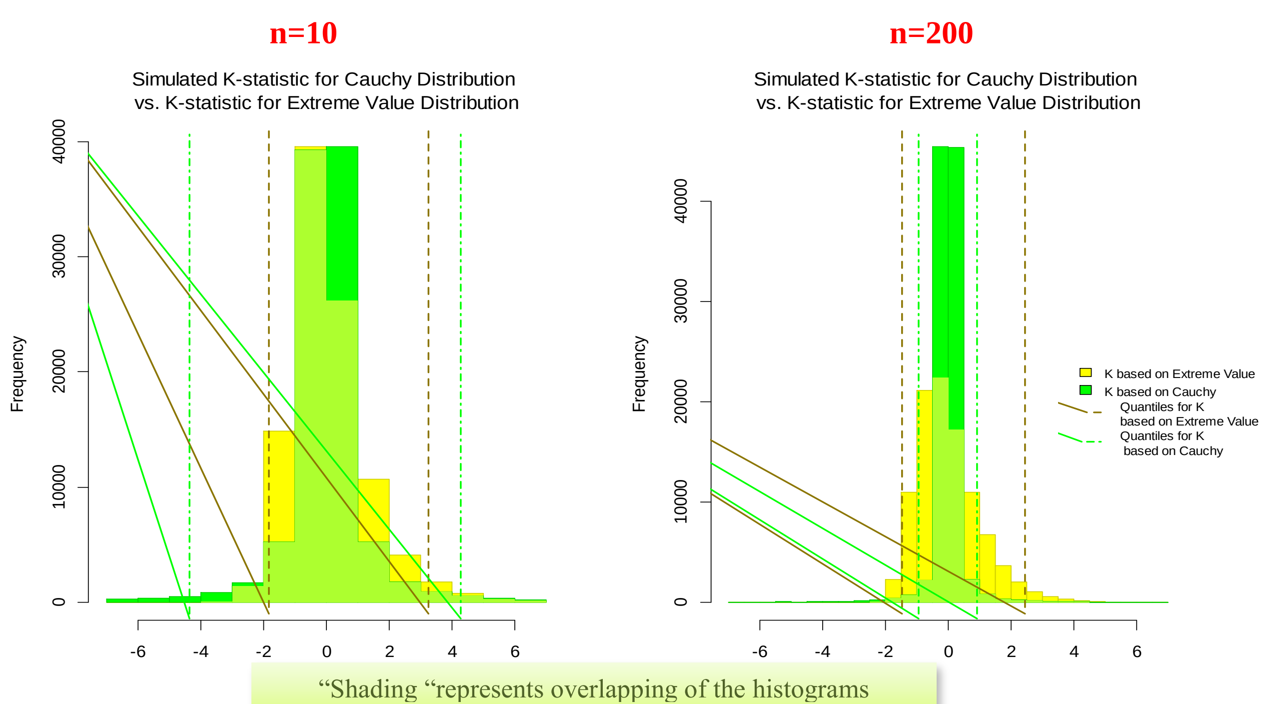

Figure 4.3 Relative Frequency Histogram of Sample Standard Deviation from Cauchy Data for small and large sample sizes (M=100000)

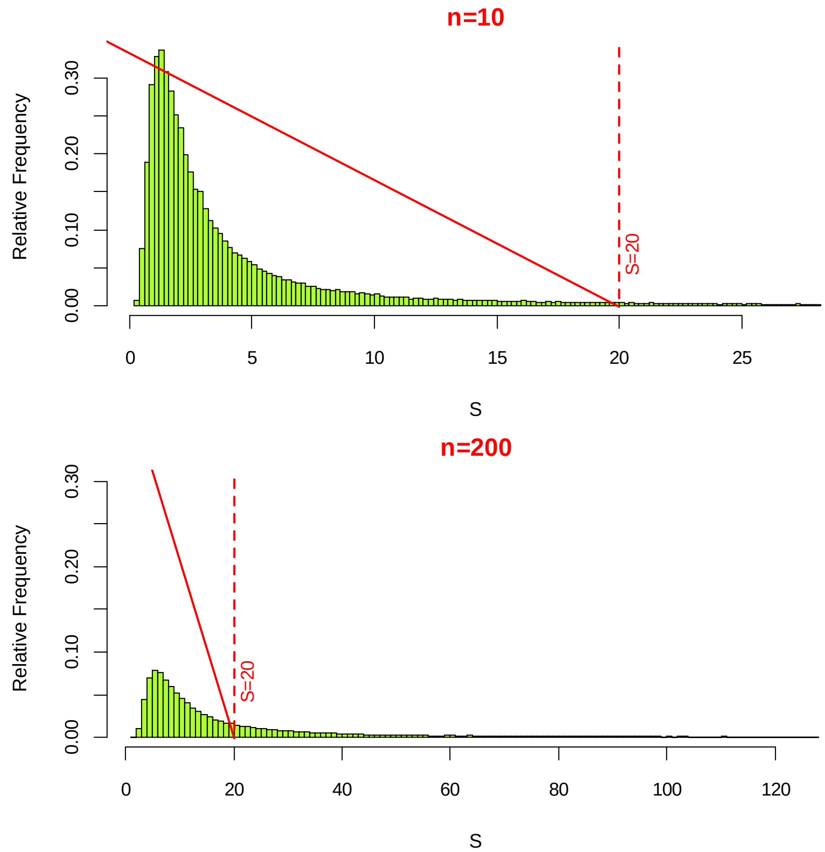

Assumed Model: Extreme Value Distribution

Table 4.3 below presents a summary of the simulation results when constructing prediction intervals under the assumption of sampling from the extreme value distribution. M = 100000 (number of iterations for generating critical values) and L = 1000 (number of PI being constructed).

Table 4.3 Estimated Average Lengths and Coverage Rates Assuming Extreme Value Distribution, *M*=100000, *L* = 1000

$1-\alpha$	n	Critical Values	Average Width / Median Width / Coverage Rate		
			Normal	Cauchy	Extreme Value
0.95	5	{-2.37, 4.19}	6.13 / 6.08 /0.9410	248.25 / 11.20 / 0.8940	6.56 /6.12 / 0.9540
	10	{-1.83, 3.24}	4.98 /4.93 / 0.9460	153.63 /12.20 / 0.9160	5.36 / 5.13 / 0.9600
	50	{-1.54, 2.55}	4.08 / 4.08 / 0.9300	90.47 / 22.91 /0.9410	4.47 / 4.39 / 0.9580
	200	{-1.49, 2.44}	3.94 / 3.93 / 0.9200	212.20 / 42.03 / 0.9700	4.31 / 4.29 / 0.9400

As again expected, the coverage rate here is close to the nominal 0.95 for data actually sampled from the extreme value distribution. For Cauchy data, however, note the substantial under coverage for the small sample sizes, n=5 and 10 and slight over coverage for n = 200. For normal data, under coverage increases with increasing sample size. These over/under coverages are likely caused by the asymmetry of the extreme value distribution and the critical values obtained from it, as seen dramatically in Table 4.3, since *K* has a symmetric distribution when obtained from samples drawn from symmetric distributions, such as the normal and Cauchy.

Figure 4.4 below shows the distributions of the *K*-statistic given in (2) based on data sampled from F_0 - the extreme value and scaled *t*-distributions, for small and large sample sizes. Sampling *K* from the scaled *t*-distribution corresponds to what happens when data are sampled from the normal distribution. For n = 10 and n = 200, it is clear that distributions of *K* based on extreme value data (green) are asymmetric. Moreover, those distributions have heavier positive tail and lighter negative tail relative to the *t*-distribution (yellow). Therefore, when using the extreme value distribution model, the critical values are shifted to the right, which is the reason given above for having under or over-coverage for data actually sampled from symmetric distributions.

Figure 4.4 Superimposed Histograms of Simulated K-statistic Based on Extreme Values Data and Normal Data for Small and Large Sample Sizes (M=100000) (M=100000)

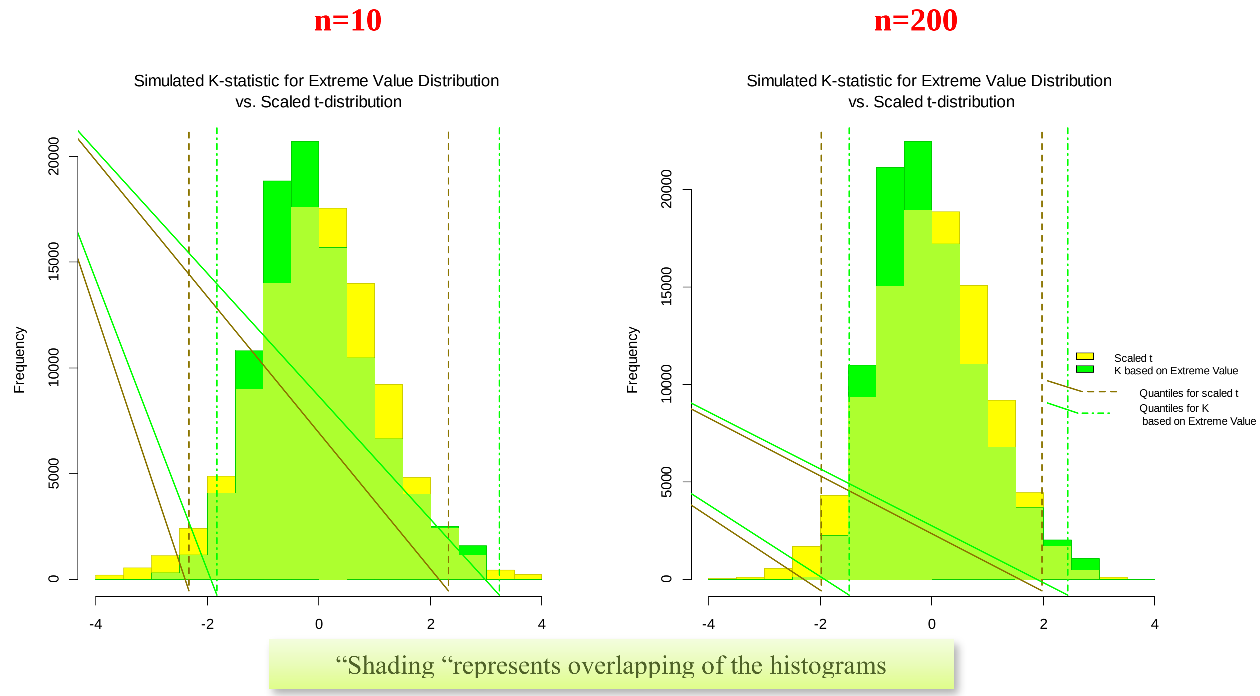

Another Look at the Simulation Study

From a practical point of view, it is often the case that a researcher has data with an unknown underlying distribution and he/she makes assumptions about the distribution in order to proceed with inferences. In this section, robustness is investigated by reconfiguring Tables 4.1-4.3 by fixing the distribution from which the data are sampled and letting the model used to construct the data vary, representing what happens when the experimenter makes the correct assumptions and makes the wrong assumptions. The tables given below summarize coverage rates and interval lengths when the data are actually sampled from the normal, Cauchy and extreme value distributions.

Normal Distribution

Suppose the unknown underlying distribution is Normal. Table 4.4 presents simulated coverage rates and 'average' widths when the data are generated from the normal distribution and, respectively, the normal, Cauchy and extreme value distribution models are used to construct the prediction intervals.

Table 4.4 Estimated Average Lengths and Coverage Rates for Normally Distributed Data

1-α	n	Average Width / Median Width / Coverage Rate		
		Normality Assumption	Cauchy Assumption	Extreme Value Assumption
0.95	5	5.71/ 5.56 /0.9460	12.45 / 12.08 / 0.9960	6.13/ 6.08 /0.9410
	10	4.53 / 4.48 / 0.9460	8.37 / 8.28 / 0.9980	4.98 /4.93 / 0.9460
	50	4.05 / 4.05 / 0.9450	3.58 / 3.57 / 0.9080	4.08 / 4.08 / 0.9300
	200	3.95 / 3.94 / 0.9610	1.76 / 1.76 / 0.6160	3.94 / 3.93 / 0.9200

As expected, the actual coverage rate that is closest to the nominal 95% rate is attained, under normality, when the experimenter's distributional assumption is correct. Among the three assumed models, the minimum interval lengths are achieved under normality for n=5, 10 and 50. However, for the large sample size n=200, average and median length of the prediction intervals are much smaller for the Cauchy model. For the Cauchy assumption, we see under coverage for small samples and over coverage for moderate and large samples, which is a lack of robustness. Note that the results obtained using the extreme value model for small sample size are very similar to those obtained using the normal model, a case where robustness with respect to one of these two possible assumed models holds.

Cauchy distribution

Table 4.5 Estimated Average Lengths and Coverage Rates for Cauchy Distributed Data

1-α	n	Average Width / Median Width / Coverage Rate		
		Normality Assumption	Cauchy Assumption	Extreme Value Assumption
0.95	5	31.47 / 10.46 / 0.8960	82.01 / 22.18 / 0.9520	248.25 / 11.20 / 0.8940
	10	56.89 / 12.18 / 0.9050	202.10 / 21.96 / 0.9480	153.63 /12.20 / 0.9160
	50	157.97 / 23.68 / 0.9550	69.75 / 20.35 / 0.9480	90.47 / 22.91 / 0.9410
	200	381.38 / 46.77 / 0.9660	175.82 / 20.73 / 0.9400	212.20 / 42.03 / 0.9700

As expected, we observe that coverage rate closest to the nominal value of 0.95 is attained when correctly assuming that the data are sampled from the Cauchy distribution. The average widths of the interval are rather unstable in most cases. Specifically, for assumptions of normality and extreme value distributions, the median width of a prediction interval increases with increasing sample size. However, for the correct assumption, the median widths of the prediction intervals stay the same for different sample sizes. To a great extent, using *PI* coverage rate as the criterion, for moderate and large sample sizes (n=50, 200), the normality and extreme value assumptions do not give significantly different result from the "correct" Cauchy assumption. However, the median width of a prediction interval is smaller for the correct Cauchy data, especially for the large sample size.

Extreme Values Distribution

Table 4.6 Estimated Average Lengths and Coverage Rates for Extreme Value Distributed Data

1-α	n	Average Width / Median Width / Coverage Rate		
		Normality Assumption	Cauchy Assumption	Extreme Value Assumption
0.95	5	6.28 / 5.72 / 0.9390	13.47 / 12.58 / 0.9910	6.56 /6.12 / 0.9540
	10	4.90 / 4.73 / 0.9400	8.87 / 8.54 / 0.9900	5.36 / 5.13 / 0.9600
	50	4.39 / 4.37 / 0.9500	3.93 / 3.88 / 0.9190	4.47 / 4.39 / 0.9580
	200	4.34 / 4.32 / 0.9570	1.94 / 1.94 / 0.6450	4.31 / 4.29/ 0.9400

Again, as expected, the actual coverage rate that is closest to the nominal 95% rate is attained using extreme value critical points, when the experimenter's assumption is correct. For the very small sample size, n = 5, assumed normality yields narrower prediction intervals and under coverage. However, for sample sizes n=10, 50, and 200 the normality assumption works even better than the correct extreme value assumption, providing coverage rates within the margin of error from the nominal 0.95 rate and the narrowest *PI* lengths. Using the Cauchy model, the widths of *PIs* decrease with increasing sample size, and become quite narrow for the large sample size, n= 200. Consequently, *PI* coverage rates drop significantly with increasing sample size. The results obtained using the normal model for large and moderate sample sizes are very similar to those obtained using the extreme value model, a case of where robustness with respect to one of two possible assumed models holds.

Chapter 5 - Discussion

Normal distribution

It was shown through simulation that the assumption of normality is only robust in some cases when applied to Cauchy and extreme value data. In terms of coverage rate, the normal model works well with normal data, extreme value data and Cauchy data with large sample size. However, the normal model does not work well when the data are Cauchy and sample sizes are small. Specifically, it was seen that for n = 5 and n = 10 the attained coverage rates is quite a bit smaller than the nominal 0.95 and *PI's* widths are much larger than from the normal distribution. For the Cauchy distribution, which has high probability of extreme values, the mean width of the *PI*'s is very unstable and actually increases with increasing sample size. As expected, median widths are much smaller and more stable than mean width for the Cauchy distribution. Mean and median widths for the normal and extreme value data are rather close, although the narrowest widths are attained for the normal distribution, a case where the assumed distribution is the "true" distribution. With increasing sample size, *PI* width decreases, which means the uncertainty about the value of a" future" observation decreases.

Cauchy Distribution

The Cauchy assumption results in a lack of robustness with respect to departures from the assumed model. As expected, the actual coverage rates are close to the nominal 0.95 for the Cauchy data since the assumption about the distribution being Cauchy is true. However, even with coverage rates are close to their nominal value, the lengths of the prediction intervals are large relative to the interquartile rang = 1.35. This happens because of the heavy tails of the Cauchy distribution and the resulting large variations in sample means and standard deviations. The median lengths for the intervals constructed from Cauchy data are more stable and much smaller than mean lengths. For normal and extreme value data, for small sample sizes, the Cauchy model prediction intervals exhibit over-coverage. For moderate and large sample sizes, the actual coverage rates are significantly less than their nominal value 0.95. As demonstrated in Figures 4.1, critical values used in constructing a *PI* computed from the Cauchy model are too small for small samples or too large for large samples when used with normal and extreme value data.

Extreme Value Distribution

The extreme value assumption exhibits a lack of robustness when used with normal and Cauchy data, largely due to the asymmetry of the distribution of the *K* statistic and the resulting critical values when sampling from the extreme value distribution. As expected, the coverage rate here is close to the nominal 0.95 for data actually sampled from the extreme value distribution. For Cauchy data, both under coverage and over coverage were observed. For normal data, one may notice under coverage for the moderate and large sample sizes. It was shown that, when using the extreme value distribution model, the *K* statistic is asymmetric and the critical values of *PI* are shifted to the right. This is one reason for having under or over coverage for data actually sampled from symmetric distributions.

Another look at the simulation

I also look at the results from a practical point of view, when it is often the case that a researcher has data with unknown underlying distribution and he/she makes assumptions about the distribution in order to proceed with inferences. In this case, robustness was investigated by fixing the distribution from which the data are sampled and letting the model used to construct the data vary, representing what happens when the experimenter makes the correct assumptions and makes the wrong assumptions. I summarized coverage rates and interval lengths when the data are actually sampled from the normal, Cauchy and extreme value distributions and found out that if data actually come from the extreme value and Cauchy distributions, for large and moderate sample sizes, both normal or extreme value assumptions work well in terms of coverage rate of *PI*. This robustness does not hold for small sample sizes. If data are actually sampled the the from normal distribution, the *PI*’s obtained using the extreme value model for small sample size are very similar to those obtained using the normal model. However, the median width of a prediction interval is smaller with the correct assumptions, especially for the large sample size.

Conclusion

In sum, my report consists of a simulation study of the robustness of prediction intervals with respect to departures from the assumed distribution. I used the representative families of location-scale distributions. I use the sample mean as an estimate of location and sample standard deviation as a scale estimate. To study robustness I evaluated and compared coverage rates, mean and median lengths of *PI*s for cases when the distribution assumption is correct versus when it is not.

It was shown through the simulation that the assumption of normality is quite robust to the departures from normality. The normal model works well in terms of the coverage rate for normal data, extreme value data and Cauchy data with large sample size. The Cauchy assumption shows lack of robustness to the departures of assumed model due to high probability of extreme values and unstable mean and standard deviation. The extreme value assumption shows a lack of robustness to the departures of assumed model due to asymmetry of the distribution.

As for a future work, one may consider to study the robustness of prediction intervals with other measurements of center and spread of the distributions. For example, the median and interquartile range could be used instead of the mean and standard deviation in forming K. Those measurements will be more stable with respect of extreme observations and will possibly give different results regarding the robustness of PI’s.

References

Abraham, Bovas and Ledolter, Johannes (1983). Statistical methods for forecasting. New York : Wiley.

Aitchison, J. and I. R. Dunsmore (1975). Statistical prediction analysis. Cambridge University Press

Baker G. A. (1935). The probability that the mean of a second sample will differ from the mean of the first sample by less than a certain multiple of standard deviation of the first sample. http://projecteuclid.org/DPubS/Repository/1.0/Disseminate?view=body&id=pdf_1&handle=euclid.aoms/1177732565

Berkowitz , Jeremy. Testing Density Forecasts, with Applications to Risk Management. Journal of Business & Economic Statistics , Vol. 19, No. 4 (Oct., 2001), pp. 465-474. http://www.jstor.org/stable/info/1392281

Chatfield C. (1993). Calculating Interval Forecasts. Journal of Business & Economic Statistics , Vol. 11, No. 2 (Apr., 1993), pp. 121-135. http://www.jstor.org/stable/10.2307/1391361

Christoffersen Peter F. (1998) Evaluating Interval Forecasts. International Economic Review , Vol. 39, No. 4, Symposium on Forecasting and Empirical Methods in Macroeconomics and Finance (Nov., 1998), pp. 841-862 . http://www.jstor.org/stable/2527341

Fisher Amy, Horn Paul S., (1994) Robust prediction intervals in a regression setting, Computational Statistics & Data Analysis, Volume 17, Issue 2, February 1994, Pages 129-140.

Geisser, Seymour (1993). Predictive inference : an introduction. New York : Chapman & Hall

Heskes Tom. Practical Confidence and Prediction Intervals (1997). Advances in Neural Information Processing Systems 9

Isengildina Olga, Irwin , Scott H. and Good Darrel L.. Evaluation of USDA Interval Forecasts of Corn and Soybean Prices. American Journal of Agricultural Economics , Vol. 86, No. 4 (Nov., 2004), pp. 990-1004. http://www.jstor.org/stable/4492787

Landon Joshua , Singpurwalla Nozer D. Choosing a Coverage Probability for Prediction Intervals. The American Statistician, May 2008, Vol. 62, No. 2. http://www.gwu.edu/~stat/irra/Irra_files/published%20papers/Prediction%20Invervals.pdf

Moore, D. S., and McCabe, G. P. (1999), Introduction to the Practice of Statistics (4th ed.), New York, NY: Freeman.

Mykland, Aslak. Financial Options and Statistical Prediction Intervals. The Annals of Statistics , Vol. 31, No. 5 (Oct., 2003), pp. 1413-1438. http://www.jstor.org

Neter, J., Kutner, M. H., Nachtsheim, C. J., and Wasserman, W. (1996), Applied Linear Statistical Models (4th ed.), Chicago, IL: Irwin.

Olive David J (2003). Prediction Intervals in the Presence of Outliers. Publisher: Southern Illinois University, Pages: 1-8. http://www.math.siu.edu/olive/pppi.pdf

Olive David J.(2007) Prediction Intervals for Regression Models. Computational Statistics and Data Analysis, Vol. 51. http://opensiuc.lib.siu.edu/cgi/viewcontent.cgi?article=1013&context=math_articles

Patel, J. K. (1989). Prediction intervals - a review. Communications in Statististics. Theory and Methods, 18(7), pp. 2393-2465

Preston Scott (2000). Teaching Prediction Intervals. Journal of Statistics Education v.8, n.3. http://www.amstat.org/publications/jse/secure/v8n3/preston.cfm

Ralph D. Snyder, J. Keith Ord and Anne B. Koehler. Prediction Intervals for ARIMA Models. Journal of Business & Economic Statistics , Vol. 19, No. 2 (Apr., 2001), pp. 217-225. http://www.jstor.org/stable/1392165

Appendix A - Additional Tables and Figures

Table A.1 Results Regarding Sensitivity of Critical Values of K statistics from Cauchy Distribution

n=10

M (time)	Critical Values
1000	-5.061481 4.238534
3000	-4.020733 4.036417
5000	-5.666179 4.151743
7000	-4.766980 4.697485
9000	-3.814126 3.898761
11000	-4.204481 4.272888
13000	-4.246926 4.474927
15000	-4.251938 4.100810
17000	-4.376495 4.333627
19000	-4.148104 4.286055
21000	-4.468018 4.131385
23000	-4.505702 4.117492
25000	-4.323016 4.151127
30000	-4.390403 4.219974
35000	-4.507743 4.151119
50000(5 sec)	-4.439949 4.595045
100000(10 sec)	-4.262407 4.324377
200000(20 sec)	-4.333312 4.316925
500000(40 sec)	-4.415021 4.339141

n=200

M (time)	Critical Values
1000	-1.3419367 1.1386466
3000	-0.9979260 0.8087007
5000	-0.8353134 0.9874517
7000	-0.9189608 0.9582606
9000	-0.9439111 0.8825460
11000	-0.8234814 0.8487343
13000	-0.9620640 0.9993206
15000	-0.8731867 0.9423853
17000	-0.8751119 0.9627143
19000	-0.8462095 0.9222872
21000	-0.9065414 0.8930322
23000	-0.8828203 0.9551173
25000	-0.8707986 0.8970807
30000	-0.8907397 0.8631322
35000(8 sec)	-0.9160072 0.9291775
50000(16 sec)	-0.9220668 0.9075826
100000 (96 sec)	-0.8946382 0.8729634
200000(7min 12 sec)	-0.9116870 0.9340537
500000	-

Figure A.1 Estimated Coverage Rates and 5% Confidence Intervals. Cauchy Assumptions

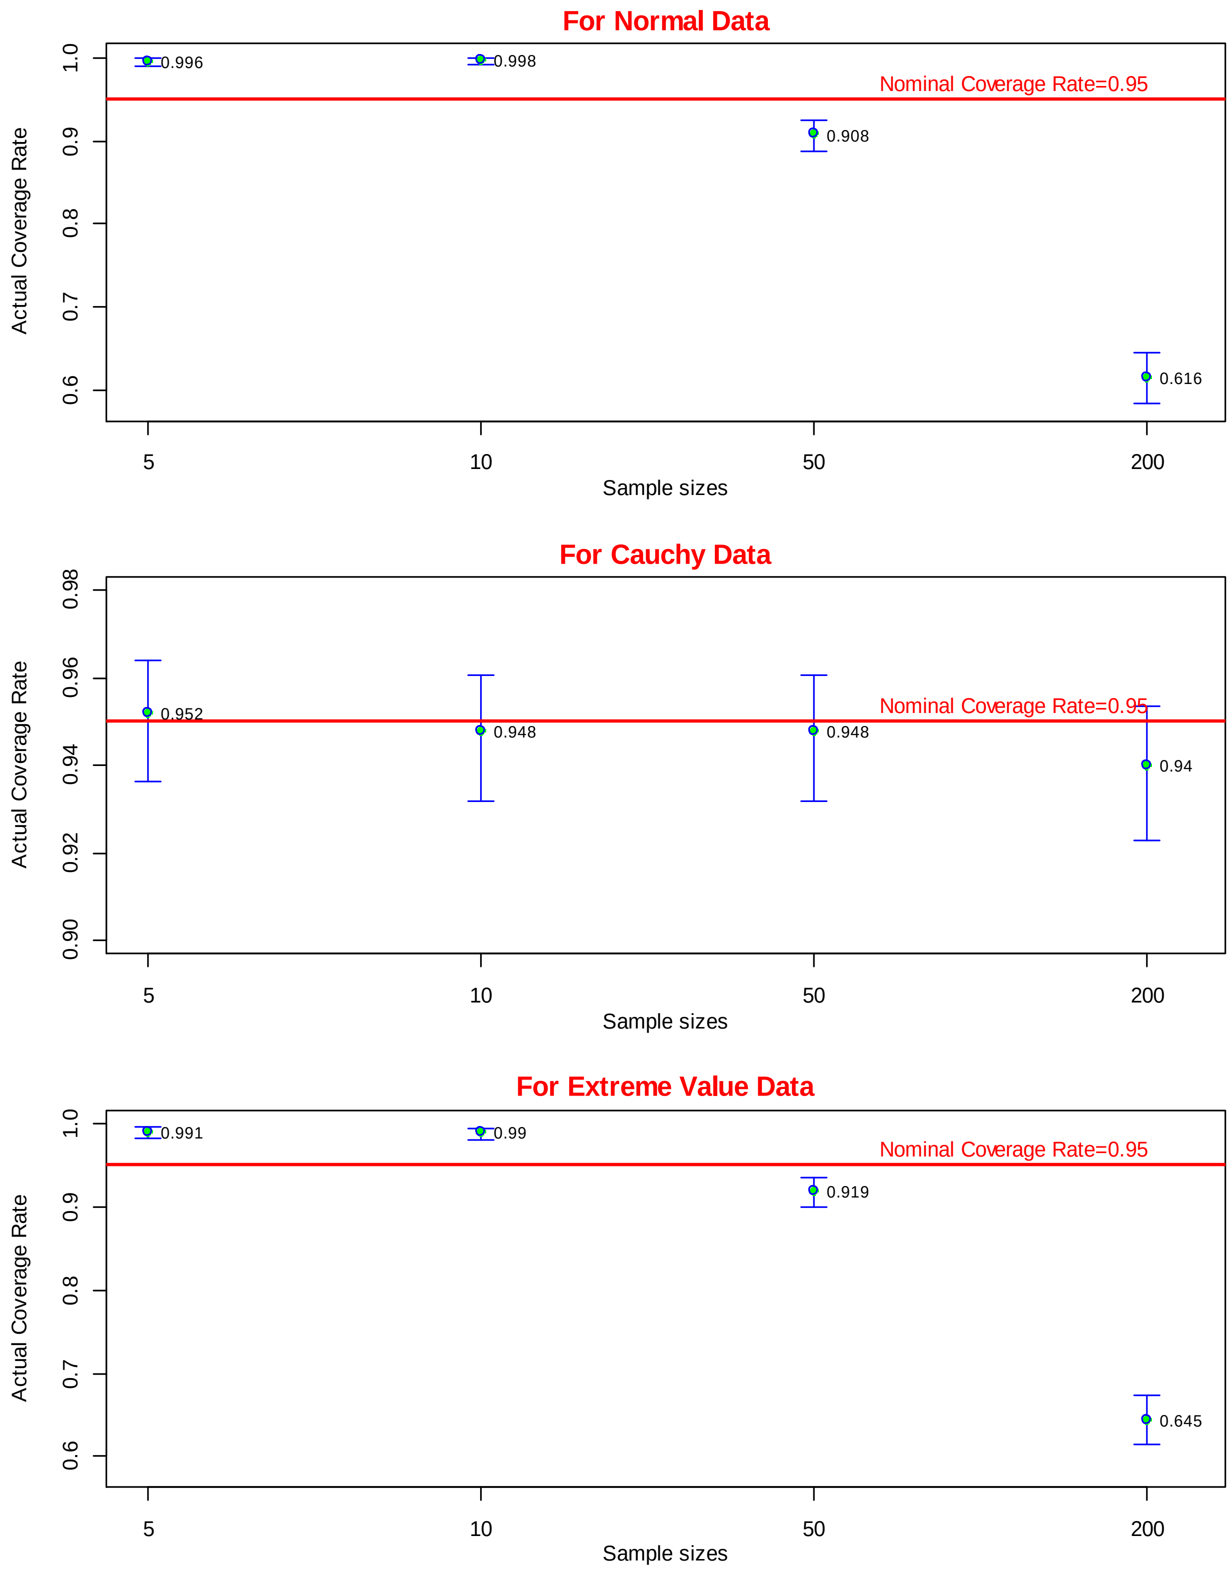

Appendix B - R Code for the Simulation Study

```
n=10 #sample size
n
M=100000 #number of iteration for generating each pair of critical values
M
L=1000 #number of PI being constructed
L
library(evd)

#______________________________________________________________________________
#Transperant Histohram
plotOverlappingHist <- function(a, b, colors=c("yellow","green","greenyellow"),
                                breaks=NULL, xlim=NULL, ylim=NULL){
  ahist=NULL
  bhist=NULL
   if(!(is.null(breaks))){
    ahist=hist(a,breaks=breaks,plot=F)
    bhist=hist(b,breaks=breaks,plot=F)
  } else {
    ahist=hist(a,plot=F)
    bhist=hist(b,plot=F)

    dist = ahist$breaks[2]-ahist$breaks[1]
    breaks = seq(min(ahist$breaks,bhist$breaks),max(ahist$breaks,bhist$breaks),dist)

    ahist=hist(a,breaks=breaks,plot=F)
    bhist=hist(b,breaks=breaks,plot=F)
  }

  if(is.null(xlim)){
    xlim = c(min(ahist$breaks,bhist$breaks),max(ahist$breaks,bhist$breaks))
  }

  if(is.null(ylim)){
    ylim = c(0,max(ahist$counts,bhist$counts))
  }

  overlap = ahist
  for(i in 1:length(overlap$counts)){
    if(ahist$counts[i] > 0 & bhist$counts[i] > 0){
      overlap$counts[i] = min(ahist$counts[i],bhist$counts[i])
    } else {
      overlap$counts[i] = 0
    }
  }
```

```
  plot(ahist,  xlim=xlim, ylim=ylim, col=colors[1], border="yellow3", main=" ", xlab=" ")
  plot(bhist,  xlim=xlim, ylim=ylim, col=colors[2], add=T, border="green3")
  plot(overlap, xlim=xlim, ylim=ylim, col=colors[3], add=T, border="transparent" )
}



#____________________________________________________________________________
#Standard Normal Assumption

#Estimate critical values for PI by simulation
Xbar<-1:M
S<-1:M
T<-1:M
for (i in 1:M) {
Z=rnorm(n+1, 0,1 )
Xbar[i]=mean(Z[1:n])
S[i]=sd(Z[1:n])
T[i]=(Z[n+1]-Xbar[i])/S[i]
}

# for n=5
lq=-2.78
uq=2.78

# for n=10
lq=-2.26
uq=2.26

# for n=50
lq=-2.01
uq=2.01

# for n=200
lq=-1.97
uq=1.97

q=c(lq,uq)
q



#Generate new data and construct PI
LL<-1:L
UL<-1:L
```

```
A<-1:L
LLc<-1:L
ULc<-1:L
B<-1:L
LLe<-1:L
ULe<-1:L
D<-1:L

set.seed(2000)

for (k in 1:L) {

# Data coming from Normal Distr
Zm=rnorm(n+1, 0, 1)
Xbarm=mean(Zm[1:n])
Sm=sd(Zm[1:n])
LL[k]=Xbarm+lq*Sm*sqrt((n+1)/n)
UL[k]=Xbarm+uq*Sm*sqrt((n+1)/n)
#Estimate if Y is inside the PI
if (Zm[n+1]<=UL[k]&&Zm[n+1]>=LL[k]) A[k]=1 else A[k]=0

# Data coming from Caushy Distr
C=rcauchy(n+1, location = 0, scale = 1.39/2)
Xbarc=mean(C[1:n])
Sc=sd(C[1:n])
LLc[k]=Xbarc+lq*Sc*sqrt((n+1)/n)
ULc[k]=Xbarc+uq*Sc*sqrt((n+1)/n)
#Estimate if Y is inside the PI
if (C[n+1]<=ULc[k]&&C[n+1]>=LLc[k]) B[k]=1 else B[k]=0

# Data coming from Extreme Value Distr
E=rgev(n+1, loc=0, scale=(1.35/-log(log(0.75)/log(0.25))), shape=0)
Xbare=mean(E[1:n])
Se=sd(E[1:n])
LLe[k]=Xbare+lq*Se*sqrt((n+1)/n)
ULe[k]=Xbare+uq*Se*sqrt((n+1)/n)
#Estimate if Y is inside the PI
if (E[n+1]<=ULe[k]&&E[n+1]>=LLe[k]) D[k]=1 else D[k]=0

}


# Simulated mean/median widths of PI
# for data coming from Normal Distr
Widthn=UL-LL
# for data coming from Caushy Distr
```

```
WidthC=ULc-LLc
# for data coming from Extreme Value Distr
WidthE=ULe-LLe

# Simulated coverage rate
# for data coming from Normal Distr
CRn=sum(A)/L
# for data coming from Caushy Distr
CRc=sum(B)/L
# for data coming from Extreme Value Distr
CRe=sum(D)/L

#Print the results for Normal Assumption
c(mean(Widthn),median(Widthn),CRn)
c(mean(WidthC),median(WidthC),CRc)
c(mean(WidthE),median(WidthE),CRe)




#___________________________________________________________________________
###Estimate critical values with Caushy DISTR

#Sensitiviity of Critical Values of T-Statistics

M<-1:13
M[1]=1000
M
q2=matrix(ncol=2, nrow=13)
for (s in 1:13) {
Xbar2<-1:M[s]
S2<-1:M[s]
T2<-1:M[s]
for (i in 1:M[s]) {
Z2=rcauchy(n+1, location = 0, scale = 1.39/2)
Xbar2[i]=mean(Z2[1:n])
S2[i]=sd(Z2[1:n])
T2[i]=(Z2[n+1]-Xbar2[i])/S2[i]
}
win.graph()
par(mfrow=c(1,2))
hist(Z2, main=expression("Simulated Cauchy Distribution"), xlab="X", font.main=1, cex=0.5)
abline(v=c(quantile(Z2, 0.025), quantile(Z2, 0.975)), lty=4)
lq2=quantile(T2, 0.025)
uq2=quantile(T2, 0.975)
q2[s,1]=lq2
```

```
q2[s,2]=uq2
hist(T2, main=c("Simulated T-statistic \n for Cauchy Distribution"), xlab="T", font.main=1,
cex=0.5)
abline(v=q2[s,], lty=4)
M[s+1]=M[s]+2000
M[s+1]
}
cbind(M,q2)

#Estimate critical values for PI
set.seed(2010)
Xbar2<-1:M
S2<-1:M
T2<-1:M
for (i in 1:M) {
Z2=rcauchy(n+1, location = 0, scale = 1.39/2)
Xbar2[i]=mean(Z2[1:n])
S2[i]=sd(Z2[1:n])
T2[i]=(Z2[n+1]-Xbar2[i])/S2[i]
}

lq2=quantile(T2, 0.025)
uq2=quantile(T2, 0.975)
q2=c(lq2,uq2)

#Histogram of S
lq22=quantile(S2, 0)
uq22=quantile(S2, 0.95)
q22=c(lq22,uq22)
S22<-S2[S2<=uq22]

win.graph()
par(mfrow=c(2,1))
par(mar=c(4,4,2,1))
#dev.set(2)
hist(S22, breaks=100,  main=" n=200 ", xlab="S", ylab="Relative Frequency", xlim=q22,
col="greenyellow", freq=FALSE, font.main=2, col.main="red", cex.axis=0.8, cex.lab=0.9)
abline(v=20, lty=2, col="red", lwd=2)
text(20, 0.04, "S=20", pos=4, col="red", srt=90, cex=0.8)


#Comparing K-statistic Based on Cauchy Distribution to T-statistic under normality
T22<-T2[T2>=-7&T2<=7] #Cut off long tails
set.seed(2011)
t=rt(M,df=n-1)
par(mfrow=c(1,2))
```

```
hist(t, breaks=100, xlim=c(quantile(t, 0.025), quantile(t, 0.975)), ylim=c(0, 33000),
main=c("Simulated K-statistic for Cauchy Distribution vs. t-distribution"), xlab="K",
font.main=1, cex=0.5, col="green")
abline(v=c(quantile(t, 0.025), quantile(t, 0.975)), lty=4, col="green")
hist(T22, add=T, col="yellow" )
abline(v=q2, lty=4, col="goldenrod3")
q2
c(quantile(t, 0.025), quantile(t, 0.975))

#Another way to make a picture
T22<-T2[T2>=-7&T2<=7] #Cut off long tails
set.seed(2011)
t=rt(M,df=n-1)
Scaled_t=t*(sqrt((n+1)/n))
Scaled_t2<-Scaled_t[Scaled_t>=-10&Scaled_t<=10] #Cut off long tails
win.graph()
#dev.set(2)

par(fig=c(0,0.44,0,1), new=F)
#par(fig=c(0.49,0.93,0,1), new=T)
par(mar=c(2,4,4,0))
plotOverlappingHist(Scaled_t,T22)
title(main="Simulated K-statistic for Cauchy Distribution
vs. Scaled t-distribution", xlab="K", font.main=1, cex=0.5)
abline(v=c(quantile(Scaled_t, 0.025), quantile(Scaled_t, 0.975)), lty=2, col="gold4",lwd=2)
abline(v=q2, lty=4, col="green", lwd=2)

par(fig=c(0.83,1,0,1), new=T)
par(mar=c(0,0,0,0))
plot(0:1, 0:1,type="n", axes=FALSE, ann=FALSE)
legend(x=0, y=0.5,c("Scaled t", "K based on Cauchy"), cex=0.75,
fill=c("yellow","green"),bty="n" )
legend(x=0, y=0.45,c("Quantiles for scaled t", "Quantiles for K
 based on Cauchy"), lty=c(2,4), cex=0.75, col=c("gold4", "green"), bty="n" , lwd=2)



#Another way to make a picture for Extreme K vs. Cauchy K
set.seed(2001)
for (i in 1:M) {
Z3=rgev(n+1, loc=0, scale=(1.35/-log(log(0.75)/log(0.25))),  shape=0)
Xbar3[i]=mean(Z3[1:n])
S3[i]=sd(Z3[1:n])
T3[i]=(Z3[n+1]-Xbar3[i])/S3[i]
}
lq3=quantile(T3, 0.025)
```

```
uq3=quantile(T3, 0.975)
q3=c(lq3,uq3)

T33<-T3[T3>=-7&T3<=7] #Cut off long tails
T22<-T2[T2>=-7&T2<=7] #Cut off long tails

win.graph()
#dev.set(2)
n
par(fig=c(0,0.44,0,1), new=F)
#par(fig=c(0.49,0.93,0,1), new=T)
par(mar=c(2,4,4,0))
plotOverlappingHist(T33,T22)
title(main="Simulated K-statistic for Cauchy Distribution
vs. K-statistic for Extreme Value Distribution", xlab="K", font.main=1, cex=0.5)
abline(v=q3, lty=2, col="gold4",lwd=2)
abline(v=q2, lty=4, col="green", lwd=2)

par(fig=c(0.83,1,0,1), new=T)
par(mar=c(0,0,0,0))
plot(0:1, 0:1,type="n", axes=FALSE, ann=FALSE)
legend(x=0, y=0.5,c("K based on Extreme Value", "K based on Cauchy"), cex=0.75,
fill=c("yellow","green"),bty="n" )
legend(x=0, y=0.45,c("Quantiles for K
based on Extreme Value", "Quantiles for K
 based on Cauchy"), lty=c(2,4), cex=0.75, col=c("gold4", "green"), bty="n" , lwd=2)




#Genetate new data and construct PI
LL2<-1:L
UL2<-1:L
A2<-1:L
LLc2<-1:L
ULc2<-1:L
B2<-1:L
LLe2<-1:L
ULe2<-1:L
D2<-1:L

for (k in 1:L) {

# Data coming from Normal Distr
Zm=rnorm(n+1, 0, 1)
```

```
Xbarm=mean(Zm[1:n])
Sm=sd(Zm[1:n])
LL2[k]=Xbarm+lq2*Sm
UL2[k]=Xbarm+uq2*Sm
#Estimate if Y is inside the PI
if (Zm[n+1]<=UL2[k]&&Zm[n+1]>=LL2[k]) A2[k]=1 else A2[k]=0
A2[k]

# Data coming from Caushy Distr
C=rcauchy(n+1, location = 0, scale = 1.39/2)
Xbarc=mean(C[1:n])
Sc=sd(C[1:n])
LLc2[k]=Xbarc+lq2*Sc
ULc2[k]=Xbarc+uq2*Sc
#Estimate if Y is inside the PI
if (C[n+1]<=ULc2[k]&&C[n+1]>=LLc2[k]) B2[k]=1 else B2[k]=0

# Data coming from Extreme Value Distr
E=rgev(n+1, loc=0, scale=(1.35/-log(log(0.75)/log(0.25))), shape=0)
Xbare=mean(E[1:n])
Se=sd(E[1:n])
LLe2[k]=Xbare+lq2*Se
ULe2[k]=Xbare+uq2*Se
#Estimate if Y is inside the PI
if (E[n+1]<=ULe2[k]&&E[n+1]>=LLe2[k]) D2[k]=1 else D2[k]=0

}


# Simulated critical values
lq2
uq2

# Simulated PI
# for data coming from Normal Distr
LL2
UL2
A2
# for data coming from Caushy Distr
LLc2
ULc2
B2
# for data coming from Extreme Value Distr
LLe2
ULe2
D2
```

```
# Simulated mean/median widths of PI
# for data coming from Normal Distr
Widthn2=UL2-LL2
mean(Widthn2)
median (Widthn2)
# for data coming from Caushy Distr
WidthC2=ULc2-LLc2
mean(WidthC2)
median(WidthC2)
# for data coming from Extreme Value Distr
WidthE2=ULe2-LLe2
mean(WidthE2)
median(WidthE2)

# Simulated coverage rate
# for data coming from Normal Distr
CRn2=sum(A2)/L
CRn2

# for data coming from Caushy Distr
CRc2=sum(B2)/L
CRc2
# for data coming from Extreme Value Distr
CRe2=sum(D2)/L
CRe2

c(mean(Widthn2),median(Widthn2),CRn2)
c(mean(WidthC2),median(WidthC2),CRc2)
c(mean(WidthE2),median(WidthE2),CRe2)

#------------------------------------------------------------------------------------
###Estimate critical values with Extreme Value DISTR
n=10
#Estimate critical values for PI
Xbar3<-1:M
S3<-1:M
T3<-1:M

set.seed(2001)
for (i in 1:M) {
Z3=rgev(n+1, loc=0, scale=(1.35/-log(log(0.75)/log(0.25))),  shape=0)
Xbar3[i]=mean(Z3[1:n])
S3[i]=sd(Z3[1:n])
T3[i]=(Z3[n+1]-Xbar3[i])/S3[i]
}
```

```
lq3=quantile(T3, 0.025)
uq3=quantile(T3, 0.975)
q3=c(lq3,uq3)
q3

# Compare Simulated T-statistic for Extreme Value Distribution vs. t-distribution
T33<-T3[T3>=-3&T3<=3] #Cut off long tails
set.seed(2011)
t=rt(M,df=n-1)
Scaled_t=t*(sqrt((n+1)/n))
Scaled_t<-Scaled_t[Scaled_t>=-4&Scaled_t<=4] #Cut off long tails
win.graph()
#dev.set(2)

par(fig=c(0,0.44,0,1), new=T)
#par(fig=c(0.49,0.93,0,1), new=F)
par(mar=c(2,4,4,0))
plotOverlappingHist(Scaled_t,T33)
title(main="Simulated K-statistic for Extreme Value Distribution
vs. Scaled t-distribution", xlab="K", font.main=1, cex=0.5)
abline(v=c(quantile(Scaled_t, 0.025), quantile(Scaled_t, 0.975)), lty=2, col="gold4",lwd=2)
abline(v=q3, lty=4, col="green", lwd=2)

par(fig=c(0.83,1,0,1), new=T)
par(mar=c(0,0,0,0))
plot(0:1, 0:1,type="n", axes=FALSE, ann=FALSE)
legend(x=0, y=0.5,c("Scaled t", "K based on Extreme Value"), cex=0.75,
fill=c("yellow","green"),bty="n" )
legend(x=0, y=0.45,c("Quantiles for scaled t", "Quantiles for K
 based on Extreme Value"), lty=c(2,4), cex=0.75, col=c("gold4", "green"), bty="n" , lwd=2)




#Genetate new data and construct PI
LL3<-1:L
UL3<-1:L
A3<-1:L
LLc3<-1:L
ULc3<-1:L
B3<-1:L
LLe3<-1:L
ULe3<-1:L
D3<-1:L
```

```
set.seed(2003)
for (k in 1:L) {

# Data coming from Normal Distr
Zm=rnorm(n+1, 0, 1)
Xbarm=mean(Zm[1:n])
Sm=sd(Zm[1:n])
LL3[k]=Xbarm+lq3*Sm
UL3[k]=Xbarm+uq3*Sm
#Estimate if Y is inside the PI
if (Zm[n+1]<=UL3[k]&&Zm[n+1]>=LL3[k]) A3[k]=1 else A3[k]=0
A3[k]

# Data coming from Caushy Distr
C=rcauchy(n+1, location = 0, scale = 1.35/2)
Xbarc=mean(C[1:n])
Sc=sd(C[1:n])
LLc3[k]=Xbarc+lq3*Sc
ULc3[k]=Xbarc+uq3*Sc
#Estimate if Y is inside the PI
if (C[n+1]<=ULc3[k]&&C[n+1]>=LLc3[k]) B3[k]=1 else B3[k]=0

# Data coming from Extreme Value Distr
E=rgev(n+1, loc=0, scale=(1.35/-log(log(0.75)/log(0.25))), shape=0)
Xbare=mean(E[1:n])
Se=sd(E[1:n])
LLe3[k]=Xbare+lq3*Se
ULe3[k]=Xbare+uq3*Se
#Estimate if Y is inside the PI
if (E[n+1]<=ULe3[k]&&E[n+1]>=LLe3[k]) D3[k]=1 else D3[k]=0


}

# Simulated critical values
lq3
uq3

# Simulated PI
# for data coming from Normal Distr
LL3[1:5]
UL3[1:5]
A3[1:5]
# for data coming from Caushy Distr
LLc3[1:5]
ULc3[1:5]
```

```
B3[1:5]
# for data coming from Extreme Value Distr
LLe3[1:5]
ULe3[1:5]
D3[1:5]

# Simulated mean/median widths of PI
# for data coming from Normal Distr
Widthn3=UL3-LL3
mean(Widthn3)
median (Widthn3)
# for data coming from Caushy Distr
WidthC3=ULc3-LLc3
mean(WidthC3)
median(WidthC3)
# for data coming from Extreme Value Distr
WidthE3=ULe3-LLe3
mean(WidthE3)
median(WidthE3)

# Simulated coverage rate
# for data coming from Normal Distr
CRn3=sum(A3)/L
CRn3
# for data coming from Caushy Distr
CRc3=sum(B3)/L
CRc3
# for data coming from Extreme Value Distr
CRe3=sum(D3)/L
CRe3

c(mean(Widthn3),median(Widthn3),CRn3)
c(mean(WidthC3),median(WidthC3),CRc3)
c(mean(WidthE3),median(WidthE3),CRe3)



#____________________________________________________________________________
#RESULTS
#Assumptoins of Normal Distr
# Simulated mean/median widths/CoverageRate of PI
c1=c(mean(Widthn),median(Widthn),CRn)
c2=c(mean(WidthC),median(WidthC),CRc)
c3=c(mean(WidthE),median(WidthE),CRe)

#Assumptoins of Caushy Distr
```

```
# Simulated critical values
# Simulated mean/median widths/CoverageRate of PI
c4=c(mean(Widthn2),median(Widthn2),CRn2)
c5=c(mean(WidthC2),median(WidthC2),CRc2)
c6=c(mean(WidthE2),median(WidthE2),CRe2)

#Assumptoins of Extreme Value Distr
# Simulated critical values
# Simulated mean/median widths/CoverageRate of PI
c7=c(mean(Widthn3),median(Widthn3),CRn3)
c8=c(mean(WidthC3),median(WidthC3),CRc3)
c9=c(mean(WidthE3),median(WidthE3),CRe3)


#______________________________________________________________________________
#Plotting Coverage rate with 95%CI
#For Cauchy Assumption
library(gplots)
win.graph()
par(mfrow=c(3,1))
ll=seq(1:4)
ul=seq(1:4)

#For Normal Data
CovRate=c(0.9960,.9980, 0.9080, 0.6160)
CovRate1=CovRate*1000
for (i in 1:4) {
ci=prop.test(x=CovRate1[i], n=L)
ll[i]=ci$conf.int[1]
ul[i]=ci$conf.int[2]
}
l=c(1, 1.5, 2, 2.5)
#win.graph()
par(mar=c(4,4,2,1))
plotCI(l,CovRate, li=ll, ui=ul, ylim=c(.58,1), xlim=c(1,2.555), main="For Normal Data", xlab="
", ylab="Actual Coverage Rate", xaxt="n",
 cex.main=1.3, cex.axis=1, cex.lab=1, pch=21, pt.bg="green", col="blue", font.lab=2,gap=0)
axis(side=1, at=l, labels=c("5", "10", "50", "200") )
mtext("Sample sizes", side=1, cex=0.8, line=2)
abline(h=0.95, col="red", lwd=2)
text(2.3, 0.96,"Nominal Coverage Rate=0.95", col="red", cex=1)
text(l,CovRate, c(0.996,.998, 0.908, 0.616), pos=4, cex=0.8)


#For Cauchy Data
CovRate=c(0.9520,0.9480, 0.9480, 0.9400)
```

```
CovRate1=CovRate*1000
for (i in 1:4) {
ci=prop.test(x=CovRate1[i], n=L)
ll[i]=ci$conf.int[1]
ul[i]=ci$conf.int[2]
}
l=c(1, 1.5, 2, 2.5)
#win.graph()
par(mar=c(4,4,2,1))
plotCI(l,CovRate, li=ll, ui=ul, ylim=c(.90,0.98), xlim=c(1,2.555), main="For Cauchy Data",
xlab=" ", ylab="Actual Coverage Rate", xaxt="n",
 cex.main=1.3, cex.axis=1, cex.lab=1, pch=21, pt.bg="green", col="blue", font.lab=2, gap=0)
axis(side=1, at=l, labels=c("5", "10", "50", "200") )
mtext("Sample sizes", side=1, cex=0.8, line=2)
abline(h=0.95, col="red", lwd=2)
text(2.3, 0.953,"Nominal Coverage Rate=0.95", col="red", cex=1)
text(l,CovRate, c(0.952,0.948, 0.948, 0.940), pos=4, cex=0.8)


#For Extreme Value Data
CovRate=c(0.9910,0.99000, 0.9190, 0.6450)
CovRate1=CovRate*1000
for (i in 1:4) {
ci=prop.test(x=CovRate1[i], n=L)
ll[i]=ci$conf.int[1]
ul[i]=ci$conf.int[2]
}
l=c(1, 1.5, 2, 2.5)
#win.graph()
par(mar=c(4,4,2,1))
plotCI(l,CovRate, li=ll, ui=ul, ylim=c(.58,1), xlim=c(1,2.555), main="For Extreme Value Data",
ylab="Actual Coverage Rate", xlab=" ", xaxt="n",
 cex.main=1.3, cex.axis=1, cex.lab=1, pch=21, pt.bg="green", col="blue", font.lab=2, gap=0)
axis(side=1, at=l, labels=c("5", "10", "50", "200"))
mtext("Sample sizes", side=1, cex=0.8, line=2)
abline(h=0.95, col="red", lwd=2)
text(2.3, 0.96,"Nominal Coverage Rate=0.95", col="red", cex=1)
text(l,CovRate, c(0.9910,0.99000, 0.9190, 0.6450), pos=4, cex=0.8)
```