6-2013

Statistical Methods for Estimating the Minimum Thickness Along a Pipeline

Shiyao Liu
Iowa State University

William Q. Meeker
Iowa State University, wqmeeker@iastate.edu

Statistical Methods for Estimating the Minimum Thickness Along a Pipeline

Abstract

Pipeline integrity is important because leaks can result in serious economic or environmental losses. Inspection information from a sample of locations along the pipeline can be used to estimate corrosion levels. The traditional parametric model method for this problem is to estimate parameters of a specified corrosion distribution and then to use these parameters to estimate the minimum thickness in a pipeline. Inferences using this method are, however, highly sensitive to the distributional assumption. Extreme value modeling provides a more robust method of estimation if a sufficient amount of data is available. For example, the block-minima method produces a more robust method to estimate the minimum thickness in a pipeline. To use the block- minima method, however, one must carefully choose the size of the blocks to be used in the analysis. In this paper we use simulation to compare the properties of different models for estimating minimum pipeline thickness, investigate the effect of using different size blocks, and illustrate the methods using pipeline inspection data.

Keywords

block minima, extreme value, maximum likelihood, simulation

Disciplines

Statistics and Probability

Comments

This preprint was published as Shiyao Liu and William Meeker, "Statistical Methods for Estimating the Minimum Thickness Along a Pipeline" *Technometrics* (2014): 164-179, doi: 10.1080/00401706.2014.915236.

Statistical Methods for Estimating the Minimum Thickness Along a Pipeline

Shiyao Liu
Department of Statistics
Iowa State University
lsy641@iastate.edu

William Q. Meeker
Department of Statistics
Iowa State University
wqmeeker@iastate.edu

January 25, 2013

Abstract

Pipeline integrity is important because leaks can result in serious economic or environmental losses. Inspection information from a sample of locations along the pipeline can be used to estimate corrosion levels. The traditional parametric model method for this problem is to estimate parameters of a specified corrosion distribution and then to use these parameters to estimate the minimum thickness in a pipeline. Inferences using this method are, however, highly sensitive to the distributional assumption. Extreme value modeling provides a more robust method of estimation if a sufficient amount of data is available. For example, the block-minima method produces a more robust method to estimate the minimum thickness in a pipeline. To use the block-minima method, however, one must carefully choose the size of the blocks to be used in the analysis. In this paper we use simulation to compare the properties of different models for estimating minimum pipeline thickness, investigate the effect of using different size blocks, and illustrate the methods using pipeline inspection data.

Key word: Block minima; Extreme value; Maximum likelihood; Simulation.

1 Introduction

1.1 Motivation and Purpose

Energy companies use pipelines to transfer oil, gas and other materials from one place to another. Manufactures of chemical products use pipelines within and between their plants. When the thickness at a location falls below a fixed threshold, there is risk of leakage that could result in serious economic loss, personal injury, or damage to the environment. It is possible to use statistical methods to estimate the minimum thickness of a pipeline. The traditional parametric statistical method of modeling the minimum is to estimate the distribution of thickness from the measured pipeline thickness data and then calculate the corresponding probability distribution of the minimum thickness. However, we usually have uncertain knowledge about the particular parent distribution that appropriately describes the data generating process. Small discrepancies in the specified parent distribution can lead to substantial bias in estimating the minimum distribution.

Extreme value theory, originating with Fisher and Tippett (1928), serves as an alternative approach to model extrema. Instead of estimating the parent distribution from observations, we accept the fact that the parent distribution is unknown. An immediate consequence of the Extreme Value Theorem is that under mild conditions, the limiting distribution of properly standardized minima (or maxima) extreme values have a generalized extreme value distribution. This distribution includes three classes of extreme value distributions as special cases and these are called the Gumbel, Fréchet, and Weibull distributions respectively. The choice among these three distributions depends on the domain of attraction of the relevant tail of the parent distribution. Here we explore the use of extreme value distributions to model minimum pipeline thickness. We use simulation to investigate the alternative procedures for estimating a minimum and apply the methods to inspection data from a three-phase pipeline (i.e., a pipeline carrying a mixture of oil, gas, and water). Our results show that whether one fits a generalized extreme value distribution or one of the special extreme value distributions under an assumed domain of attraction has a large effect on the choice of block size.

1.2 Pipeline Data

In some pipeline integrity applications it is possible to do in-line pipeline inspection (ILI) by using a "smart-pig" utilizing magnetic flux leakage or ultrasonic testing technology to detect and measure corrosion and other metal-loss features in a pipeline. The smart-pig is pulled through the pipeline acquiring information with high spatial resolution (e.g., 3mm).

In most pipeline applications, however, such in-line inspections are impossible and the pipeline operators must rely on external inspections that are done at a set of sample locations, known as Thickness Measurement Locations (TMLs). The resulting sample data are then used to make inferences about the integrity of the entire pipeline. Ultrasonic and radiographic (X-ray) testing are the most commonly-used external inspection methods to measure pipeline thickness.

To illustrate the application of the different methods that one can use to estimate the distribution of a minimum in applications like pipeline integrity, we use data from a three-phase pipeline that had an original thickness of 0.375 inches. To protect sensitive information, the name of company that provided the data and the location of the pipeline cannot be disclosed.

The raw pipeline data that we received were ILI pipeline inspection data giving the location (in feet, measured from one end of the pipeline), size, and depth of observed metal-loss features. We partitioned the data into features (a feature is an indication of metal loss) observed in the 32,272 one-foot segments along the pipeline. The smart-pig identified features in 5,649 of these one-foot segments. In the other 26,623 segments, there was no detectable metal loss. Within the 5,649 locations with detectable metal loss, the number of features that were recorded ranges between 1 and 27, as shown in Figure 1. For each one-foot segment with observed metal loss, we take the minimum of all of the thickness measurements as the wall thickness response within that one-foot segment.

Analysis of the data suggests, in agreement with knowledge that we have gained from experts in pipeline industry, that metal loss tends to concentrate in certain areas of the pipeline. These points of concentration tend to be in the area of certain physical characteristics of the pipeline such as near supports, places where the slope of the pipeline changes, and near

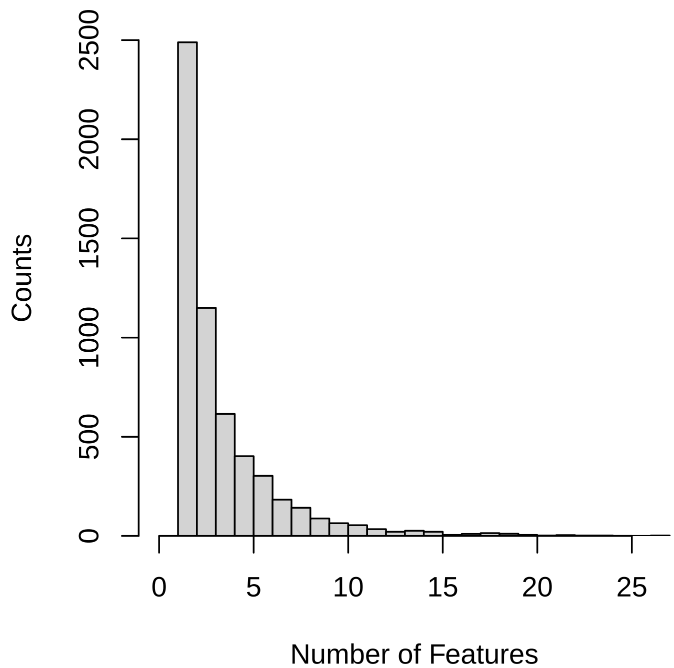

Figure 1: Number of metal-loss features recorded among the 5,649 one-foot segments that had one or more detected features.

welds that join two sections of pipe. Thus we will take as the population the 5,649 locations with detectable metal loss. To simulate the common kind of external inspection, we draw simple random samples of size $n = 200$ (approximately the smallest sample size required to use the more robust statistical methods based on extreme-value theory) and $n = 1{,}000$ (approximately the largest sample size we have seen in external inspection pipeline applications) of the one-foot segments with metal loss. Such selective sampling would correspond, roughly, to the industry practice of over-sampling at locations in a pipeline where one would expect to see higher rates of metal loss.

Because of measurement resolution limitations of the smart-pig system, (resulting from an analog-to-digital conversion with a limited number of bits per reading to allow storage of large amounts of data) the pipeline wall measurements are not known exactly due to round-off error. As suggested by Vardeman and Lee (2005), we therefore treated the data as interval-censored observations.

1.3 Some Previous Work on Extreme Value Analysis

Extreme value analysis has been used widely in many areas of application ranging from insurance and finance to meteorology and hydrology. There is a large number of books and articles regarding both the mathematical theory and applications of extreme value analysis. Gumbel (1958) is one of the earliest books and is still an important reference in extreme value analysis. Coles (2001) describes the common approaches of extreme value analysis including the block maxima (or minima) method and the threshold excess models. Castillo, Hadi, Balakrishnan and Sarabia (2005) focus particularly on applications of extreme value analysis in the engineering areas. Engeland, Hisdal and Frigessi (2004) use extreme value methods to model hydrological floods and droughts. Kowaka et al. (1994) use extreme value statistical methods to investigate corrosion phenomena. Laycock, Cottis and Scarf (1990), Laycock and Scarf (1993) and Scarf and Laycock (1994) apply the extreme value analysis to corrosion and propose a four parameter time-dependent model to extrapolate of extreme pit depths into future exposure time and larger area of metal. Shibata (1994) reviews the application of the extreme value statistics to corrosion using several examples. Scarf and Laycock (1996) use extreme value theory to model the maximum penetration caused by pitting corrosion on metal surfaces. Fougères, Holm and Rootzén (2006) design and analyze experiments to compare treatments with extreme responses, using corrosion experiments to illustrate their approach.

1.4 Overview

The remainder of this paper is organized as follows. Section 2 briefly describes the traditional statistical method of using a parametric distribution to model the minimum directly and illustrates this method with a pipeline thickness example. Section 3 introduces the extreme value distributions, presents the block minima method, and shows how to apply the block minima method with the Gumbel and the generalized extreme value distributions to estimate quantiles of the distribution of a minimum over the population. Section 4 gives the details of the design of a simulation experiment for comparing the different methods of estimating a minimum. Section 5 presents the simulation results and investigates the effect of block size choice on MSE in the block minima method. Section 6 explores the effect of block size choice on the interval estimates by comparing the relative likelihood profile plots for quantiles of the distribution of a minimum. Section 7 compares different methods of estimating quantiles of the distribution of a minimum and the corresponding confidence intervals with the pipeline thickness inspection data. Section 8 provides some summary conclusions and recommendations and suggests some areas for future research.

2 The Traditional Statistical Method to Estimate a Minimum

2.1 Methods for Estimating the Distribution of a Minimum

For independent and identically distributed (iid) random variables $X_1, X_2, \ldots, X_M$ with a cumulative distribution function (cdf) $F(x; \boldsymbol{\theta})$, the distribution of the minimum $Y_M = \min\{X_1, X_2, \ldots, X_M\}$ can be expressed as:

$$\Pr[Y_M \leq x_c] = 1 - [1 - F(x_c; \boldsymbol{\theta})]^M \tag{1}$$

where M is the population size and $\boldsymbol{\theta}$ is parameter vector. From (1), in order to estimate the minimum distribution, one needs first to specify the parent distribution $F(x; \boldsymbol{\theta})$. Then substituting an estimate of $\boldsymbol{\theta}$ provides an estimate of the distribution of the minimum.

2.2 Application of the Distribution of a Minimum

In this section, we use the pipeline wall thickness inspection data to illustrate the application of the traditional statistical method to model a minimum. In order to estimate the probability that the minimum wall thickness of the population is less than a critical limit

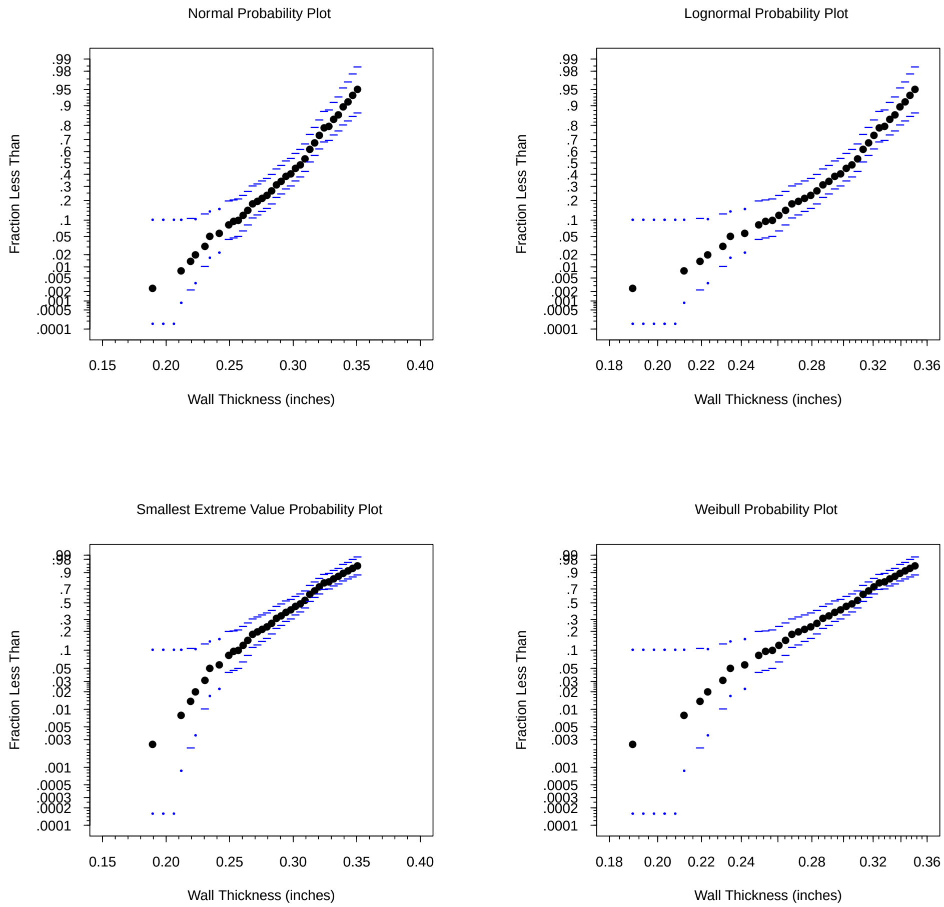

Figure 2: Probability plots for the $n = 200$ pipeline wall thickness inspection data with 95% simultaneous confidence bands.

(say, 0.10 inches in this application), one needs first to choose an appropriate parent distribution to fit the wall thickness data from the $n = 200$ locations. Because inferences on the minimum thickness generally require extrapolation into the lower tail of the distribution, the inferences can be highly sensitive to the assumed distribution. When the data are consistent with more than one distribution, it is important to do sensitivity analysis to assess the effect that different distributional assumptions will have on the final answers.

Probability plots (described, for example, in Chapter 6 of Meeker and Escobar 1998) provide a useful graphical method for assessing the adequacy of an underlying parent distribution. Figure 2 provides the normal, lognormal, smallest extreme value (SEV) and Weibull probability plots for the $n = 200$ pipeline thickness observation. These probability plots indicate a good fit for the Weibull distribution, although the lognormal distribution is also consistent with the data and provides a reasonable description, at least in the lower tail of the distribution. Therefore, we choose Weibull and lognormal distributions as candidate parent distributions to make comparative statements about the probability that the minimum thickness is less than a specified value. The Weibull and lognormal distribution cdfs are

$$F(x;\mu,\sigma) = \Phi_{\text{sev}}\left[\frac{\log(x)-\mu}{\sigma}\right] \text{ and } F(x;\mu,\sigma) = \Phi_{\text{nor}}\left[\frac{\log(x)-\mu}{\sigma}\right].$$

Here $\Phi_{\text{sev}}(z) = 1 - \exp[-\exp(z)]$ is the standardized ($\mu = 0$, $\sigma = 1$) smallest extreme value (or SEV) cdf and $\Phi_{\text{nor}}(\cdot)$ is the standard ($\mu = 0$, $\sigma = 1$) normal cdf. Although one might question the assumption of independence in this application to estimate the distribution of a minimum, if the dependence is positive, as might be expected, the value given by the method is conservative (e.g., Chapter 2 of Barlow and Proschan 1975).

Here, we illustrate the traditional method to estimate a minimum (i.e., the minimum method). For the Weibull distribution, the ML estimate of the probability that the minimum pipeline wall thickness out of the $M = 5{,}649$ one-foot segments (i.e., the population) is less than 0.10 would be:

$$\begin{aligned} \Pr[\min(X_1, X_2, \ldots, X_{5649}) \le 0.10] &= 1 - [1 - \Pr(X_1 \le 0.10)]^{5649} \\ &= 1 - \left[1 - \Phi_{\text{sev}}\left(\frac{\log(0.10) + 1.157}{0.089}\right)\right]^{5649} \\ &= 0.0144. \end{aligned} \tag{2}$$

Here $\widehat{\mu} = -1.157$ and $\widehat{\sigma} = 0.089$ are respectively the maximum likelihood estimates of the SEV parameters based on the logarithm of the wall-thickness measurements from a simple random sample of 200 randomly chosen one-foot segments out of the 5,649 one-foot segments in the pipeline. Similarly, for the lognormal distribution,

$$\Pr[\min(X_1, X_2, \ldots, X_{5649}) \le 0.10] = 1 - \left[1 - \Phi_{\text{nor}}\left(\frac{\log(0.10) + 1.2089}{0.1148}\right)\right]^{5649} \approx 0, \tag{3}$$

where $\widehat{\mu} = -1.2089$ and $\widehat{\sigma} = 0.1148$ are respectively the maximum likelihood estimates of the mean and standard deviation of log thickness based on the same sample of $n = 200$ out of the $M = 5{,}649$ pipeline wall thickness measurements.

As expected, the Weibull distribution is more conservative than the lognormal distribution in terms of estimating the probability that the minimum thickness is less than a critical limit.

3 Methods for Estimating a Minimum Based on Extreme Value Theory

3.1 Extreme Value Distributions

As mentioned in Section 2, because the parent distribution function F is not always known and inferences on the minimum imply extrapolation into the lower tail of the distribution, the use of (1) to estimate distribution of minima carries risk of serious bias. Extreme value theory provides an alternative method of modeling a minimum. Classic extreme value theory gives the asymptotic distribution for a minimum (e.g., Section 9.1.1 of Castillo, Hadi, Balakrishnan and Sarabia 2005 or Section 3.2 of Coles 2001). The limiting distribution of the minima belongs to one of the three forms known as the Gumbel, Fréchet, and Weibull families (and there are corresponding distributions for maxima that we will not explicitly consider here).

The three limiting distributions are embedded in the minimum generalized extreme value (GEV) family with a cdf

$$G(x) = 1 - \exp\left\{ - \left[1 - \xi \left(\frac{x-\mu}{\sigma} \right) \right]^{-1/\xi} \right\}, \tag{4}$$

where $\xi \neq 0$ and $1 - \xi(x-\mu)/\sigma \geq 0$. The GEV family has three parameters: a location parameter, $-\infty < \mu < \infty$, a scale parameter, $\sigma > 0$, and a shape parameter, $-\infty < \xi < \infty$. The limit of (4) as $\xi \to 0$, leads to the minimum Gumbel family with cdf

$$G(x) = 1 - \exp\left[- \exp\left(\frac{x-\mu}{\sigma} \right) \right], \quad -\infty < x < \infty. \tag{5}$$

The quantiles of the GEV distribution are obtained by solving $G(x_p) = p$ for x_p giving:

$$x_p = \begin{cases} \mu + \dfrac{\sigma}{\xi}\left\{1 - \left[-\log(1-p)\right]^{-\xi}\right\}, & \text{for } \xi \neq 0 \\ \mu + \sigma \log\left[-\log(1-p)\right], & \text{for } \xi = 0. \end{cases} \tag{6}$$

If the parent distribution F has a minimum limiting distribution G, then F is said to be in the minima domain of attraction of G. In many practical applications, physical considerations will indicate the particular form of G. Table 9.5 of Castillo, Hadi, Balakrishnan and Sarabia (2005) also summarizes the maxima and minima domain of attraction of these three types of parametric limiting distributions. From that table, the normal, the SEV distribution (minimum Gumbel distribution) and the LEV distribution (maximum Gumbel distribution) all belong to the Gumbel minima domain of attraction.

Although the minimum Gumbel and the SEV distributions are equivalent, for clarity of purpose, we use the term minimum Gumbel to refer to a limiting distribution of minima and SEV to refer to a parent distribution.

3.2 The Block Minima Method

To estimate the distribution of minima from the thickness measurements using the extreme value distributions, we need to obtain data from a minimum process. The block minima method (described, for example, in Section 3.3.1 of Coles 2001) provides an alternative method of estimating the distribution of minima from the thickness measurements by grouping the data into blocks of equal or approximately equal size and taking as data the minimum in each block. For n iid observations $X_1, \ldots, X_n$, let m denote the number of blocks, so there are $B = n/m$ observations in each block. Let $X_{\min_i} = \min\{X_{B(i-1)+1}, \ldots, X_{Bi}\}, i = 1, \ldots, m$ be the minimum value in block i. Then the block minima $X_{\min_1}, \ldots, X_{\min_m}$ are independent observations that will follow, approximately, a minimum extreme value distribution (either one of the minimum-type distributions or the generalized extreme value distribution).

3.3 Estimating the Minimum over a Population

In order to estimate the minimum of a population, when using the block minima method, one will generally need to extrapolate further into the tail of the extreme value distribution that is estimated by using the block minima method (e.g., Section 1.1.3 of Glegola 2007). In particular, this additional extrapolation is needed because when using the block minima method with blocks of size B, we obtain an estimate of the parent distribution of minima for a population of size M/B. If the constructed block minima data set $\{X_{\min_1}, \dots, X_{\min_m}\}$ follows a minimum distribution $G_B(x)$ (either the Gumbel or the GEV distribution) corresponding to blocks of size B, then the minimum thickness $Y_M = \min\{X_1, \dots, X_M\}$ of the population with size M can be treated as the minimum of a sample of M/B independent block minima with blocks of size B and the minimum distribution of Y_M is

$$
\begin{aligned}
G_M(x_c) &= \Pr(Y_M \le x_c) \\
&= 1 - \Pr\left(Y_{\min_1} > x_c, \dots, Y_{\min_{\frac{M}{B}}} > x_c\right) \\
&= 1 - [1 - G_B(x_c)]^{M/B}, \qquad (7)
\end{aligned}
$$

where $Y_{\min_j} = \min\{X_{B(j-1)+1}, \dots, X_{Bj}\}, j = 1, \dots, M/B$ is the minimum value in block j (having size B) and $Y_M = \min\{Y_{\min_1}, \dots, Y_{\min_{M/B}}\}$.

If one wants to control $G_M(x)$, such that $G_M(x) = \Pr(Y_M \le x) = p$, then one would choose the threshold to be $x_p = G_M^{-1}(p)$, the p quantile x_p of the distribution of the population minimum Y_M. The translation to the adjusted quantile in terms of the block minima distribution $G_B(x)$ is as follows:

$$
x_p = G_M^{-1}(p) = G_B^{-1}\left(1 - (1-p)^{B/M}\right) = G_B^{-1}(p^*), \qquad (8)
$$

where $p^* = 1 - (1-p)^{B/M}$. It is this quantile that will be the focus of our simulation to study estimation performance in Sections 4, 5 and 6.

4 Design of the Simulation Experiment and Simulation Details

4.1 Objective of the Simulation

Cox, Isham and Northrop (2002) investigate the asymptotic variance of the quantile estimates of a distribution of maxima under different estimation methods. Asymptotic variances do not take into account the bias. Bias, however, can play a central role in evaluation of the accuracy of estimators in finite samples. Extreme value theory is based on large-sample asymptotic results. It is important to understand how methods based on this large-sample theory will perform with finite samples. The objective of this section is to describe the design of a simulation experiment and to suggest criteria with which we can compare various statistical methods for modeling and making inferences about the minimum thickness of a pipeline and other applications where the block minima method might be used. As we will see, choice of block size plays an important role in the performance of the block minima estimation method. We then use this simulation to explore the impact of block size in the block-minima extreme value method for estimating small quantiles of a distribution of a minimum.

4.2 Experimental Factors and Their Levels

In the simulation, we used three parent distributions: the normal, the SEV (minimum Gumbel) and the LEV (maximum Gumbel) distributions. All of these distributions belong to the Gumbel minima domain of attraction. In the simulation we used sample sizes $n = 200$ (approximately the smallest sample size that would be suitable when using the block minima method) and $n =$ 1,000 (approximately the largest sample size we have seen in pipeline applications). In the block minima method, with the sample size $n = 200$, the observations were divided into equal-size blocks with sizes in the set $S_{200} = \{2, 4, 10, 20, 40\}$. With sample size $n =$ 1,000, the block sizes in our evaluations were chosen from $S_{1000} = \{5, 10, 20, 50, 100\}$.

The ratio of the population size to the sample size M/n affects the effective amount of extrapolation. The three levels for this factor are taken from the set $M/n = \{10, 100, 1000\}$.

The methods used to model the distribution of a minimum in this simulation are:

- Method 1 corresponds to the traditional minimum method described in Section 2.1 where the form of the parent distribution is specified.
- Method 2 is based on the block minima method (abbreviated as BL_{min}) described in Section 3.2 using the minimum Gumbel distribution to describe the minima of the blocks. This would be the appropriate estimation method if the parent distribution is known to be one of the distributions that has the minimum Gumbel distribution as its limiting distribution (e.g., the normal, SEV, or LEV that are used as parent distributions in the simulation).
- Method 3 is also based on the block minima method, using the minimum GEV distribution to fit the block minima data. This would be an appropriate method to use if there were no information about the underlying parent distribution.

The probability plots for the pipeline wall thickness data (Figure 2 in Section 2.2) and physical knowledge about the thickness data (they must be positive and were obtained by taking the maximum value in each one-foot segment) suggest that the parent distribution of the pipeline wall thickness can be adequately described by the distribution in the Weibull minima domain of attraction (e.g., the Weibull or the lognormal distribution). Thus, the distribution of the logarithm of the pipeline wall thickness is in the Gumbel minima domain of attraction and the simulation results can be used as the guide to model the logarithm of the pipeline wall thickness.

In practice, one seldom knows the most appropriate parametric form of the parent distribution. Without knowledge of the parent distribution, one might use an inadequate parametric distribution to estimate the minimum distribution. A slight discrepancy in the parent distribution can cause a substantial bias error in the estimation of the minimum distribution. To illustrate the sensitivity to an incorrect choice of a parent distribution, in simulation, we also use different assumed parent distributions.

4.3 ML Estimation and Comparison Criteria

The maximum likelihood (ML) method is used for estimating the parameters in the extreme value distribution. For details on the ML estimation of the parameters in the two-parameter extreme value distribution, see, for example, Meeker and Escobar (1998) and Lawless (2002). Coles (2001) provides more details on the ML method for estimating the parameters in the generalized extreme value distribution.

We use the usual definitions of mean square error (MSE), variance, and bias for comparison of estimators. For an unknown quantity θ with $\widehat{\theta}$ as an estimator, the MSE of $\widehat{\theta}$ is:

$$\mathrm{MSE}\left(\widehat{\theta}\right) = \mathrm{E}\left[\left(\widehat{\theta} - \theta\right)^2\right] = \mathrm{Var}\left(\widehat{\theta}\right) + \left[\mathrm{Bias}\left(\widehat{\theta}\right)\right]^2 \tag{9}$$

where $\mathrm{Bias}\left(\widehat{\theta}\right) = \mathrm{E}\left(\widehat{\theta} - \theta\right)$.

In the simulation, we evaluate these properties of the ML estimators of the lower quantiles of the distribution of the minimum. The ML estimators of the p quantile x_p are obtained by substituting the ML estimators of the parent distribution parameters into the quantile expressions in Section 3.1.

5 Simulation Results: Effect of Block Size on MSE

We investigated plots of the MSE, variance, and bias for the three different parent distributions. The ordering of the estimation method MSE curves for sample size $n = 1{,}000$ are generally similar to those for sample size $n = 200$. Thus, we will primarily display MSE results for $n = 200$. Subsequently, we will provide other plots that help understand the variance-bias trade-off.

5.1 Graphical Summary of MSE Results for the Normal Parent Distribution

Figure 3 compares the MSEs of ML estimators of quantiles ranging from $x_{0.0015}$ to $x_{0.15}$ of the distribution of the minimum for the different combinations of number of blocks for samples of size $n = 200$ using a normal parent distribution and ratio $M/n = 10$. Figure 4 displays similar simulation results for the normal parent distribution when the sample size $n = 1{,}000$.

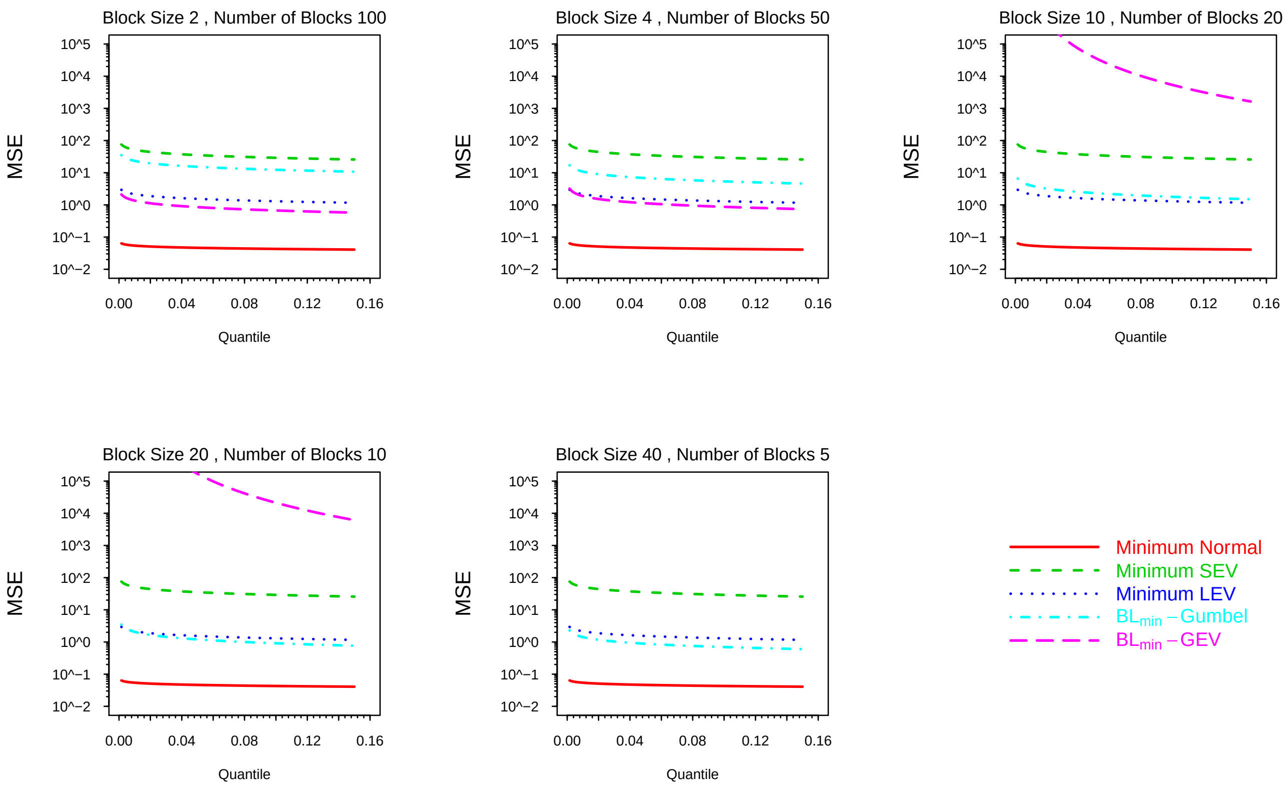

Figure 3: A comparison of the MSEs of the ML quantile estimators of the minimum for the normal parent distribution when the sample size is $n = 200$ and ratio is $M/n = 10$.

As seen in Figures 3 and 4, with data from a normal parent distribution, and small block sizes, using the $\mathrm{BL_{min}}$-Gumbel method results in ML estimators with a large MSE caused by large bias because the asymptotic extreme value approximation is poor. With a sample size $n = 200$, and ratio $M/n = 10$, a block size of at least 20 is needed for the distribution of minima to be adequately described by the $\mathrm{BL_{min}}$-Gumbel method. In further simulations (details not given here, but note Figure 6 in Section 5.3), as the ratio M/n increases from 10 to 1,000, resulting in a larger amount of effective extrapolation, bias is amplified, and the needed block size increases, say to 40.

As seen in Figure 4, with a sample size $n = 1{,}000$ and ratio $M/n = 10$, a block size

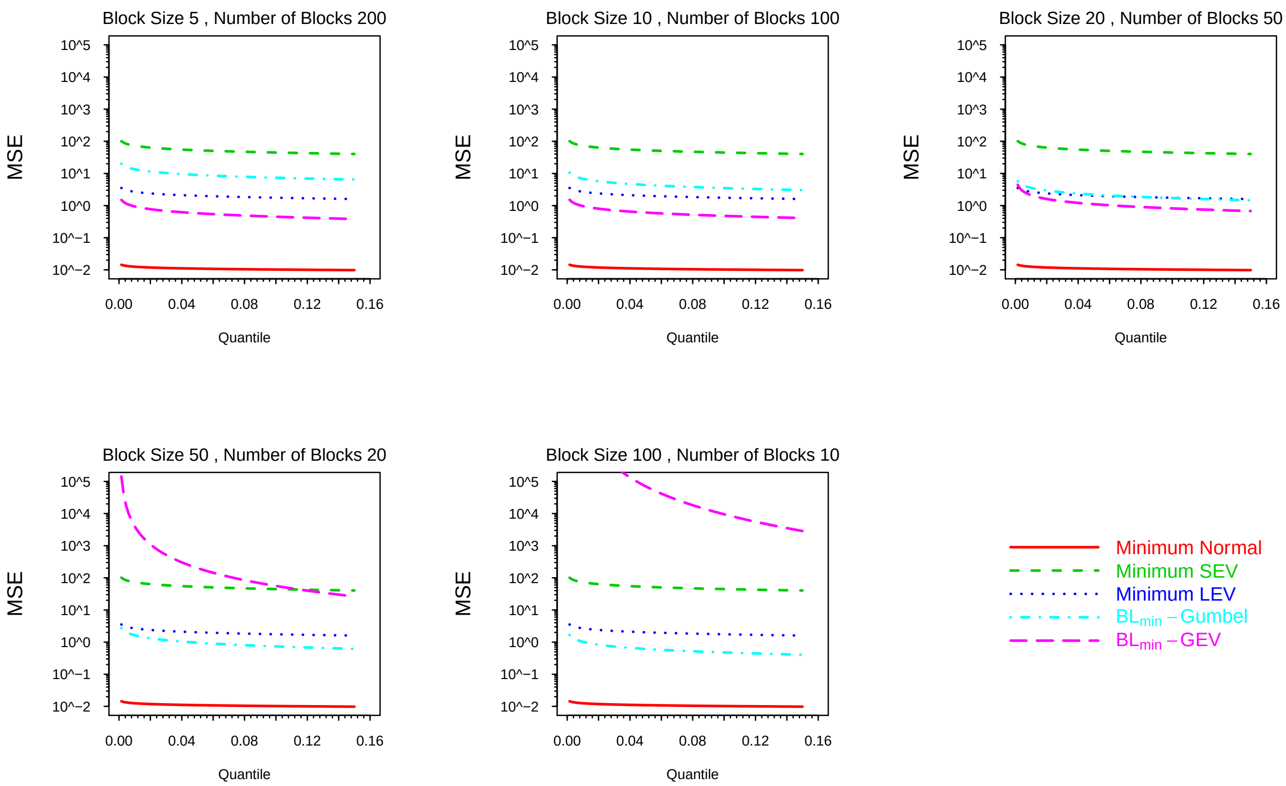

Figure 4: A comparison of the MSEs of the ML quantile estimators of the minimum for the normal parent distribution when the sample size is $n = 1{,}000$ and ratio is $M/n = 10$.

of at least 20 is again needed for the BL_{min}-Gumbel method. As the ratio M/n increases from 10 to 1,000 (again, details not given here), the block size needed to compensate for the additional bias and to provide good performance increases to say 50. The BL_{min}-GEV method, however, results in ML estimators with relatively small bias, even with small block sizes. Generally, however, the variance is much larger with the BL_{min}-GEV method because an additional shape parameter must be estimated. In Figure 3, when the number of blocks is 5, the MSE curve of the quantile estimates using the BL_{min}-GEV method is so large that it is off-scale.

For an actual normal parent distribution, the MSEs of the quantile estimators using the distribution-of-minimum method based on the mis-specified SEV parent distribution are much greater than the MSEs of quantile estimators based on the other minimum distributions. This is true for all values of the ratio M/n. The minimum-distribution estimators based on the mis-specified LEV parent distribution, however, have relatively small MSEs. This difference in behavior is because the lower tail of the LEV (SEV) distribution is similar to (different from) the lower tail of the normal distribution.

5.2 Explanation of MSE Results for the Normal Parent Distribution

Here we look in more detail at the sampling distributions of the ML estimators of quantiles of the distribution of a minimum in order to better understand the reasons for the behaviors seen in Section 5.1 and to provide insight into the choice of estimation method and block size.

The box plots in Figure 5 show the empirical sampling distribution of $\widehat{x}_{0.05}$ using the BL_{min}-Gumbel and BL_{min}-GEV methods for different block sizes $B = \{4, 10, 20\}$ when using a normal parent distribution with sample size $n = 200$ and ratio $M/n = 10$. The box plots provide some insight into the reasons for the differing behaviors (in bias and variance) between the BL_{min}-GEV and BL_{min}-Gumbel methods. In particular, when the number of blocks m is not large (say 20 or less), the BL_{min}-GEV method generates a substantial fraction of extremely small estimates, resulting in both large bias and variance. The median of the sampling distributions from the BL_{min}-GEV method, however, remains relatively close to the truth even when the block size is small. This is in contrast to the BL_{min}-Gumbel method where there is substantial bias when applied to the minimum of small blocks of normally-distributed variates.

5.3 The Effect of More Extreme Extrapolation

Figures 3, 4 and 5 provide basic comparisons for a modest amount of extrapolation into the lower tail of the distribution (i.e., $M/n = 10$). In order to compare the BL_{min}-Gumbel and BL_{min}-GEV methods with larger amounts of extrapolation, Figure 6 displays a pair of box plots of the empirical sampling distributions of $\widehat{x}_{0.05}$ using for a normal parent distribution with block sizes $B = 4$ and $B = 20$ when sample size $n = 200$ with different ratios $M/n =$

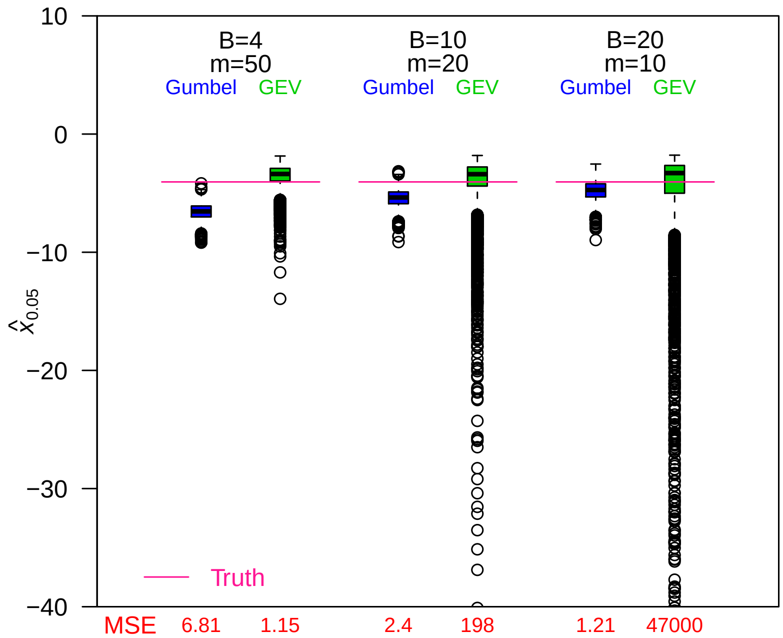

Figure 5: A comparison of the sampling distributions of $\widehat{x}_{0.05}$ using the $\mathrm{BL_{min}}$-Gumbel and $\mathrm{BL_{min}}$-GEV methods for the normal parent for $n = 200$ and $M/n = 10$ and different combinations of B and m. The horizontal lines indicate the position of the true 0.05 quantile of the distribution of the minimum in the population. Note that some $\mathrm{BL_{min}}$-GEV estimates are off scale for $m = 10$ and 20.

$\{10, 100, 1000\}$.

Figure 6 (a) shows results from samples of size $n = 200$ from a normal parent and blocks size $B = 4$. We chose $B = 4$ for this example because the MSE of the quantile estimates using the $\mathrm{BL_{min}}$-GEV method with block size $B = 4$ are the smallest among all choices of the blocks in Figure 5. The biases of the quantile estimators using the $\mathrm{BL_{min}}$-GEV method are less than the biases of quantile estimators using the $\mathrm{BL_{min}}$-Gumbel method for all ratios M/n. As the ratio increases, the biases of quantile estimators using both $\mathrm{BL_{min}}$-GEV and $\mathrm{BL_{min}}$-Gumbel methods increase. The biases of quantile estimators using the $\mathrm{BL_{min}}$-Gumbel method increase more rapidly than the biases using the $\mathrm{BL_{min}}$-GEV method. As the ratio

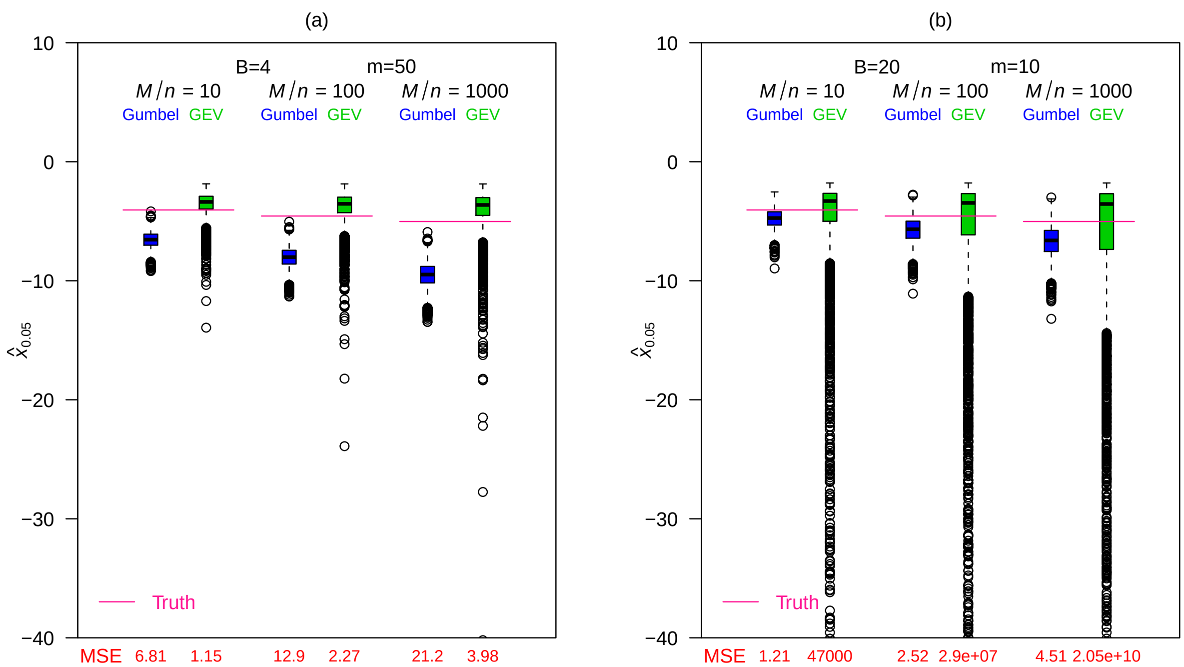

Figure 6: A comparison of the sampling distributions of $\widehat{x}_{0.05}$ using the $\mathrm{BL}_{\min}$-Gumbel and $\mathrm{BL}_{\min}$-GEV methods for the normal parent when the sample size is $n = 200$ with different combinations of M/n, B and m. The horizontal lines indicate the position of the true 0.05 quantile of the distribution of the minimum in the population. Note that some $\mathrm{BL}_{\min}$-GEV estimates are off scale for $m = 10$.

M/n becomes fairly large, say 1,000, the $\mathrm{BL}_{\min}$-GEV method, particularly when the number of the blocks (effective sample size) is not large, will generate a substantial fraction of small outliers, contributing to increased bias and variance. The behaviors of the $\mathrm{BL}_{\min}$-GEV and $\mathrm{BL}_{\min}$-Gumbel methods in the above box plots are consistent with what we observed before in the MSE plots in Figures 3 and 4. Figure 6 (b) shows box plots, similar to Figure 6 (a), but with the larger block size $B = 20$. We chose $B = 20$ for this example because the MSE of the quantile estimators using the $\mathrm{BL}_{\min}$-Gumbel method with block size $B = 20$ are the smallest among all choices of the blocks in Figure 5. We see that with the larger blocks (and thus a smaller number of blocks for estimation), the $\mathrm{BL}_{\min}$-GEV method has a large variance due to the existence of a substantial number of small outliers. Also, as the ratio M/n increases, resulting in a large amount of effective extrapolation, the variance increases tremendously and the MSE of the $\mathrm{BL}_{\min}$-GEV quantile estimators grows explosively.

5.4 MSE Results for the LEV and SEV Parent Distributions

Here we look at behavior of the competing estimation methods under alternative parent distributions. Figures 7 and 8 are similar to Figure 3 and provide comparisons of the MSEs of the quantile estimators of the distribution of the minimum in the population under the LEV and the SEV parent distributions, respectively, for a sample size $n = 200$, and a ratio $M/n = 10$.

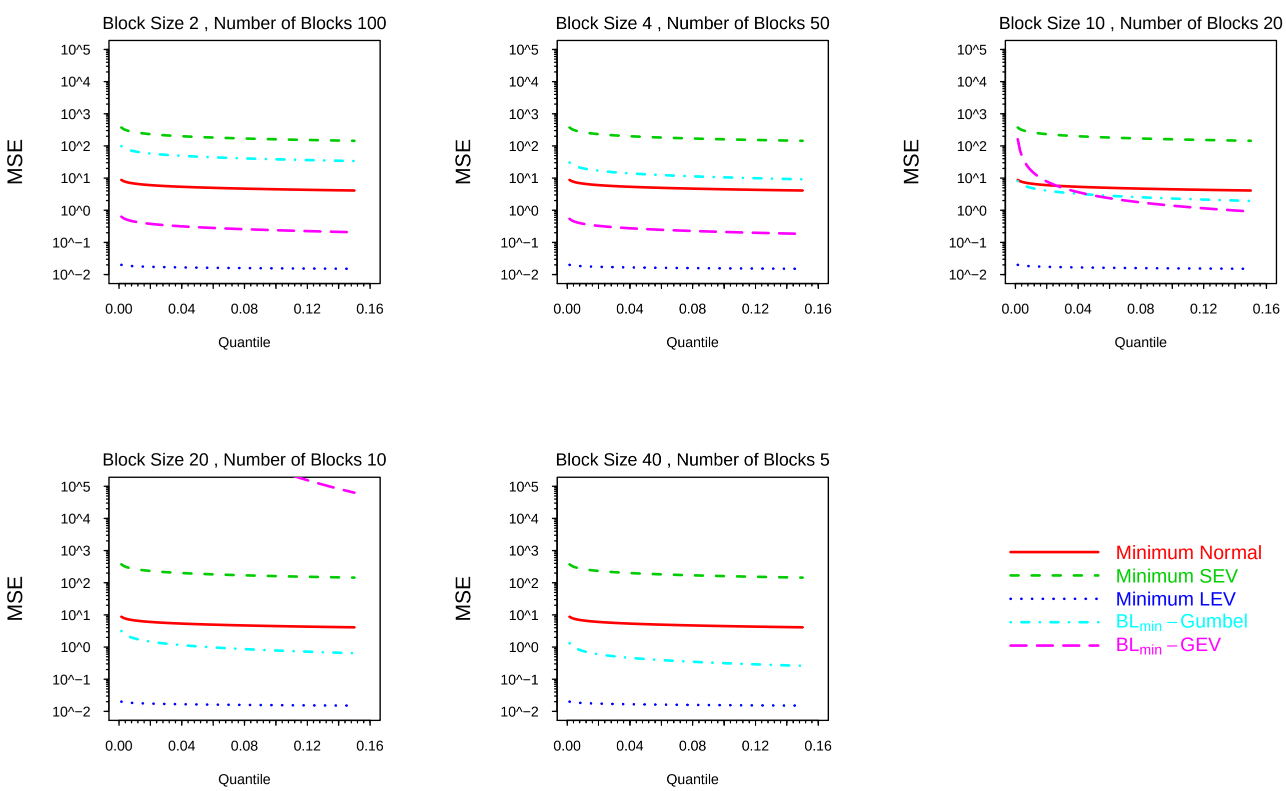

Figure 7: A comparison of the MSEs of the ML quantile estimators of the minimum for the LEV parent distribution when the sample size is $n = 200$ and ratio is $M/n = 10$.

As can be seen in Figure 7, with data from the LEV parent distribution, and small block

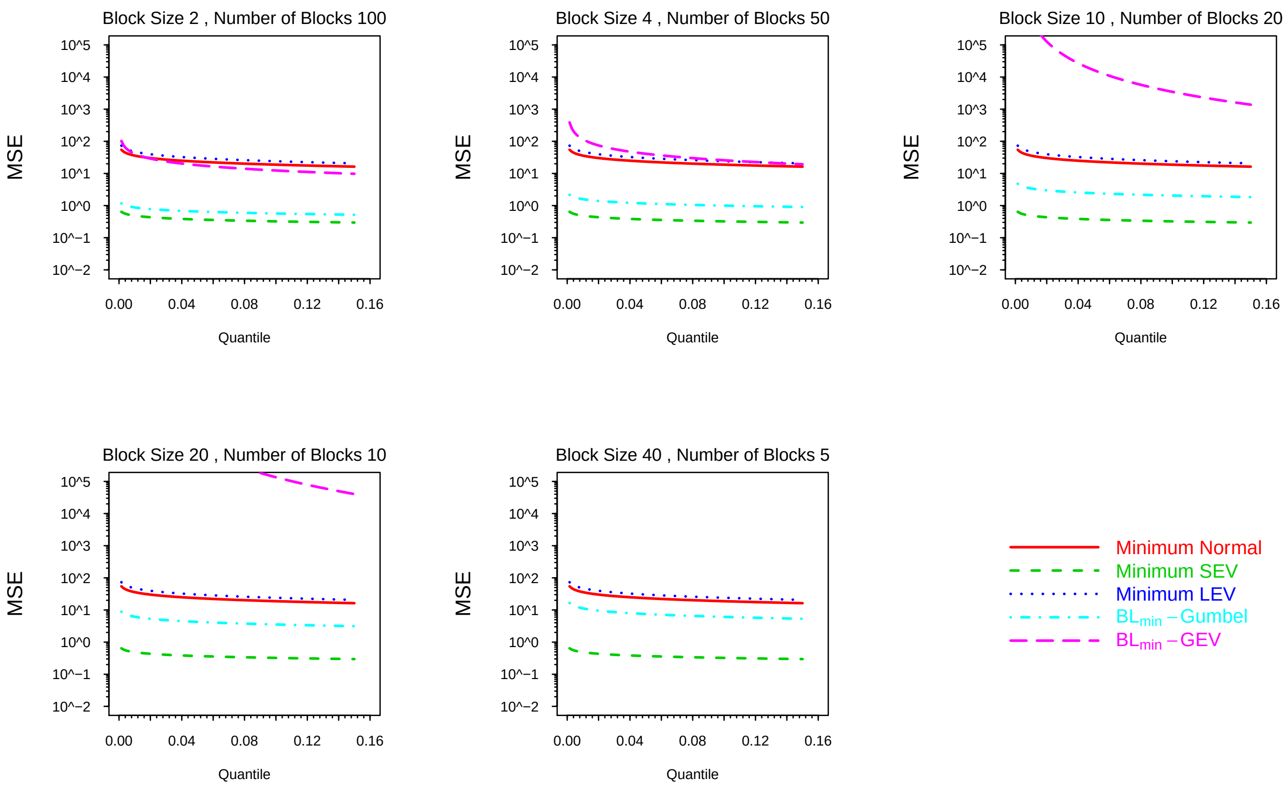

Figure 8: A comparison of the MSEs of the ML quantile estimators of the minimum for the SEV parent distribution when the sample size is $n = 200$ and ratio is $M/n = 10$.

sizes, using the BL_{min}-Gumbel method again results in ML estimators with a relatively large MSE unless the block size is large (e.g., 20 or more). We know from box plots similar to Figures 5 and 6 (not shown here) that poor performance of the BL_{min}-Gumbel method is due mostly to negative bias in the estimates of quantiles x_p with small p. On the other hand, Figure 7 also suggests that the BL_{min}-GEV method performs relatively well when there is a large number of blocks (say more than 50), even if those blocks are not large. In simulation results not displayed here, these conclusions remain the same as the ratio M/n varies from 10 to 1,000.

In Figure 7, for the LEV parent distribution, the MSEs of quantile estimators using the minimum distribution based on the mis-specified SEV parent distribution are very large. The MSE curves for the quantile estimators using the mis-specified minimum normal distribution, however, behaves much better when compared with the performance of the mis-specified minimum SEV distribution. This is because the lower tail behavior of the normal distribution is more similar to that of the LEV than it is to the SEV.

In Figure 8, for the SEV parent distribution, the distribution of the minimum is exactly the minimum Gumbel distribution, and there is no model-specification bias when using the $\mathrm{BL_{min}}$-Gumbel method. As the number of blocks increases, the MSE of the quantile estimates using the $\mathrm{BL_{min}}$-Gumbel method decreases due primarily to reduction in variance.

5.5 Impact of the Shape Parameter on the GEV Quantile Estimates

To understand why some estimates of the quantiles can be extremely small when using the $\mathrm{BL_{min}}$-GEV method, especially with a small number of blocks, we looked at scatter plots of $\mathrm{BL_{min}}$-GEV method quantile estimates versus the corresponding shape parameter estimates. Figure 9 (a), for example, shows estimates of $x_{0.05}$ for sample size $n = 200$ and block size $B = 4$ (so there are $m = 50$ blocks in the sample). The plot shows that the small estimates of the quantile $x_{0.05}$ result when the shape parameter estimates are large. The range of shape parameter estimates under the SEV parent distribution is larger than those under the normal and the LEV parent distributions, resulting in the smaller quantile estimates with the SEV parent. For the LEV and the normal parent distributions, because the sampling distribution of the shape parameter estimates does not extend far into the positive range, the bias of the $\mathrm{BL_{min}}$-GEV quantile estimators is not as large as it is for the SEV parent.

Figure 9 (b) shows a similar scatter plot for block size $B = 10$ (number of blocks $m = 20$). Compared with Figure 9 (a), the shape parameter estimates using the $\mathrm{BL_{min}}$-GEV method in Figure 9 (b) are, overall, larger and the smallest quantile estimates are much smaller than those in Figure 9 (a). The variances of the GEV parameter estimates are large when the number of blocks is small, leading to large variances for the estimators of the quantiles.

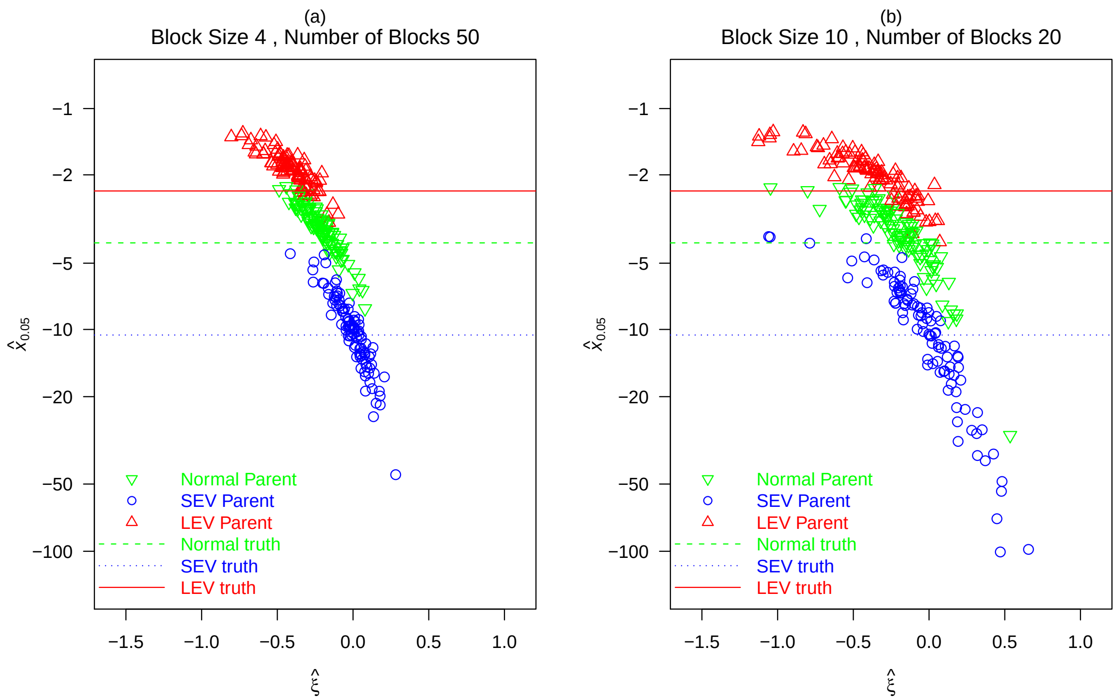

Figure 9: Scatter plot of the shape parameter estimates $\widehat{\xi}$ versus the $\widehat{x}_{0.05}$ quantile estimates for the $\mathrm{BL}_{\min}$-GEV method for normal, SEV and LEV parent distributions when sample size is $n = 200$ and ratio is $M/n = 10$. In part (a), the block size is $B = 4$ and the number of blocks is $m = 50$; In part (b), the block size is $B = 10$ and number of blocks is $m = 20$.

6 Simulation Results: Effect of Block Size and Parent Distribution on Confidence Intervals

In the previous sections, our discussion focused on investigating the properties of point estimators of the quantiles of the minimum distribution. Interval estimates that quantify the uncertainty of the point estimator are usually needed in statistical analysis. This was certainly true in our pipeline example. Although one can generally expect that confidence interval procedures based on point estimators with good (poor) properties will lead to well (poorly) behaved confidence interval procedures, in this section we compare quantile relative likelihood profile plots that give a sense of the effect that the block-size choice will have on the interval estimates. These profile plots also provide an alternative graphical tool for comparing the behaviors of the quantile estimates under different block size and distribution combinations. Because quantile relative likelihood profile curves are random, we plotted multiple realizations of these curves, corresponding to simulated data sets. In the pipeline wall thickness application, physical knowledge, the measurement process, and the data (Figure 2) suggest that the parent distribution in most cases can be adequately described by the Weibull or the lognormal distribution (i.e., the logarithm of pipeline wall thickness can be adequately described by the SEV or the normal distribution). Thus, in this section, we will mainly investigate quantile relative likelihood profile plots under the normal and the SEV parent distributions.

6.1 Simulation Results on Profile Likelihoods for the Normal Parent Distribution

Here we look at the confidence interval of the ML estimators of the quantiles of the distribution of the minimum under the normal parent distribution. Both plots in Figure 10 show 50 profile curves. The sampling distribution of the quantile point estimators can be visualized by looking at the maxima of the profiles. Additionally, the width of the likelihood-based approximate confidence intervals can be used to visually assess the precision of the corresponding point estimator.

The simulation results displayed in Figures 3 and 4 suggest that if the parent distribution is the normal distribution, a fairly large number of blocks is needed to use the $\mathrm{BL}_{\mathrm{min}}$-GEV method and a large block size is needed to use the $\mathrm{BL}_{\mathrm{min}}$-Gumbel method. As explained at the end of Section 5.4, with an SEV parent distribution, however, we know that there is no model specification bias when using the $\mathrm{BL}_{\mathrm{min}}$-GEV and $\mathrm{BL}_{\mathrm{min}}$-Gumbel methods to estimate small quantiles and the use of blocks would not be needed at all. For our likelihood profile evaluations for the $\mathrm{BL}_{\mathrm{min}}$-GEV method, we used a block size $B = 4$ (which with $n = 200$ gives 50 blocks). For the $\mathrm{BL}_{\mathrm{min}}$-Gumbel method, if we know that the parent is the normal distribution, Figure 3 suggests the block size $B = 20$. If we know parent is the SEV distribution, Figure 8 suggests the block size $B = 1$ (i.e., no blocking). Thus, we use a compromise block size $B = 10$ (giving 20 blocks).

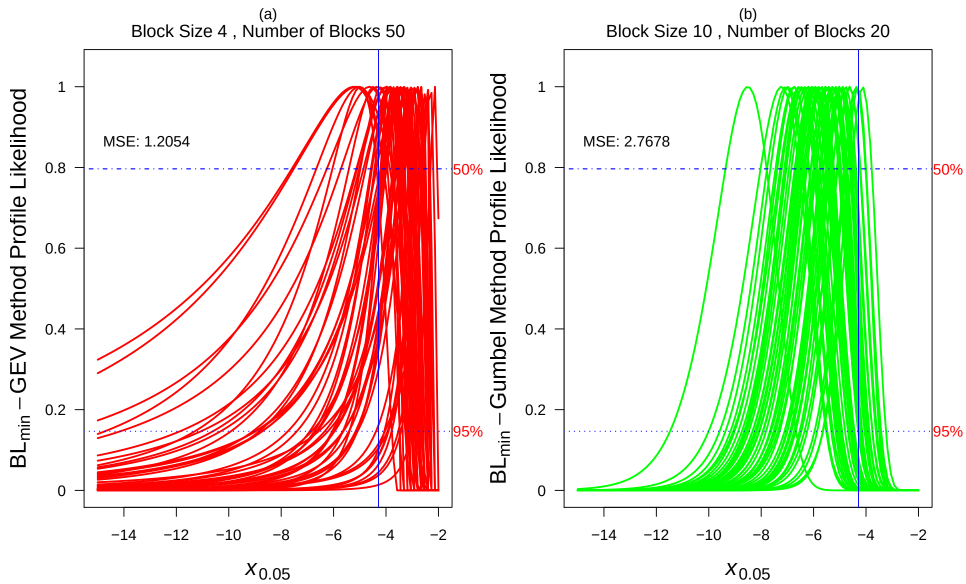

Figure 10: Comparison of quantile relative likelihood profile plots using (a) BL_{min}-GEV with the block size $B = 4$ and (b) BL_{min}-Gumbel with the block size $B = 10$ when the sample size is $n = 200$, the population size is $M = 5{,}649$, quantile is $x_{0.05}$ and the parent distribution is normal. The vertical line indicates the position of the true 0.05 quantile of the distribution of the minimum in the population. The horizontal lines allow visualization of corresponding approximate 50% and 95% likelihood-based confidence intervals.

The quantile relative likelihood profile plots in Figure 10 show that the BL_{min}-Gumbel method with a block size $B = 10$ results in less variability than the quantile relative likelihood profile plots using BL_{min}-GEV method with block size $B = 4$, especially in the lower endpoints of the likelihood-based approximate 50% and 95% confidence intervals (the horizontal lines in these plots are based on a simple chi-square distribution calibration). The MSE of the quantile estimates using the BL_{min}-GEV method, however, is smaller than that using the BL_{min}-Gumbel method. Similar simulations using $n = 1{,}000$ data sets and correspondingly larger number of blocks (details not shown here) gave similar results except that, as expected, precision was improved.

6.2 Simulation Results on Profile Likelihoods for the SEV Parent Distribution

Figure 11 provides comparisons of quantile relative likelihood profile plots using the SEV parent distribution with a sample size $n = 200$. The $\mathrm{BL_{min}}$-GEV method performs poorly, even with 50 blocks. The results in Figure 9 for the SEV parent help explain this behavior. As expected, due to the lack of model-specification bias, the MSE of the quantile estimators using the $\mathrm{BL_{min}}$-Gumbel method with block size $B = 10$ performs well.

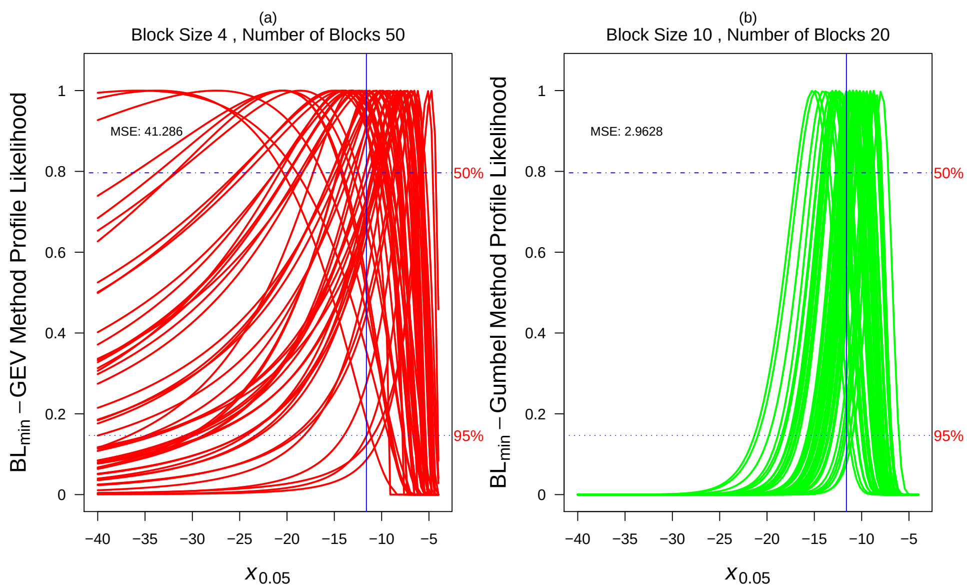

Figure 11: Comparison of quantile relative likelihood profile plots using (a) $\mathrm{BL_{min}}$-GEV with the block size $B = 4$ and (b) $\mathrm{BL_{min}}$-Gumbel with the block size $B = 10$ when the sample size is $n = 200$, the population size is $M = 5{,}649$, quantile is $x_{0.05}$ and parent distribution is the SEV distribution. The vertical line indicates the position of the true 0.05 quantile of the distribution of the minimum in the population. The horizontal lines allow visualization of corresponding approximate 50% and 95% likelihood-based confidence intervals.

6.3 General Conclusion from the Profile Likelihood Simulations

If the parent distribution is close to the lognormal distribution (normal on the log scale), our results (e.g., in Figures 3, 4 and 10) indicate that the combination of a smaller block size (e.g., $B = 4$ when n=200; $B = 10$ when $n = 1{,}000$) and the $\mathrm{BL}_{\min}$-GEV method is an appropriate choice according to the MSE criterion. If the parent distribution is close to the Weibull distribution (SEV on the log scale), the results suggest that one should use the $\mathrm{BL}_{\min}$-Gumbel method to estimate small quantiles of a distribution and, for the sake of robustness, choose a moderately large block size (e.g., $B = 10$ when n=200 or $B = 20$ when $n = 1{,}000$).

7 Estimation of the Minimum Thickness in the Pipeline

Here we return to the pipeline wall thickness inspection data and compare the estimates of $p = 0.05$ quantile and the corresponding likelihood-based approximate confidence intervals under the different methods for estimating the distribution of a minimum thickness. As in Section 2, the pipeline wall thickness data consist of two simple random samples of size $n = 200$ and $n = 1{,}000$ locations from the population of $M = 5{,}649$ measurements at locations that had metal-loss features along the three-phase pipeline.

	$n = 200$					$n = 1{,}000$				
Method	B	m	Lower	$\widehat{x}_{0.05}$	Upper	B	m	Lower	$\widehat{x}_{0.05}$	Upper
Minimum Weibull	1	200	0.0985	0.1119	0.1248	1	1000	0.1153	0.1214	0.1275
Minimum lognormal	1	200	0.1727	0.1825	0.1913	1	1000	0.1845	0.1888	0.1930
Minimum Fréchet	1	200	0.1924	0.1991	0.2050	1	1000	0.1926	0.1955	0.1981
$\mathrm{BL}_{\min}$-Gumbel	10	20	0.0965	0.1325	0.1605	20	50	0.0909	0.1123	0.1312
$\mathrm{BL}_{\min}$-GEV	4	50	0.0334	0.1452	0.1847	5	200	0.0529	0.1098	0.1464

Table 1: ML estimates and likelihood-based approximate 95% confidence intervals for $\widehat{x}_{0.05}$.

Table 1 lists the ML estimates and the corresponding lower and upper limits of the likelihood-based approximate 95% confidence intervals for the 0.05 quantile of the minimum distributions using different methods. Figure 12 displays the relative likelihood profile plots for the 0.05 quantile estimates using the block minima method. Figure 13 shows estimates of the parent and minimum distributions for the pipeline wall thickness data on Weibull

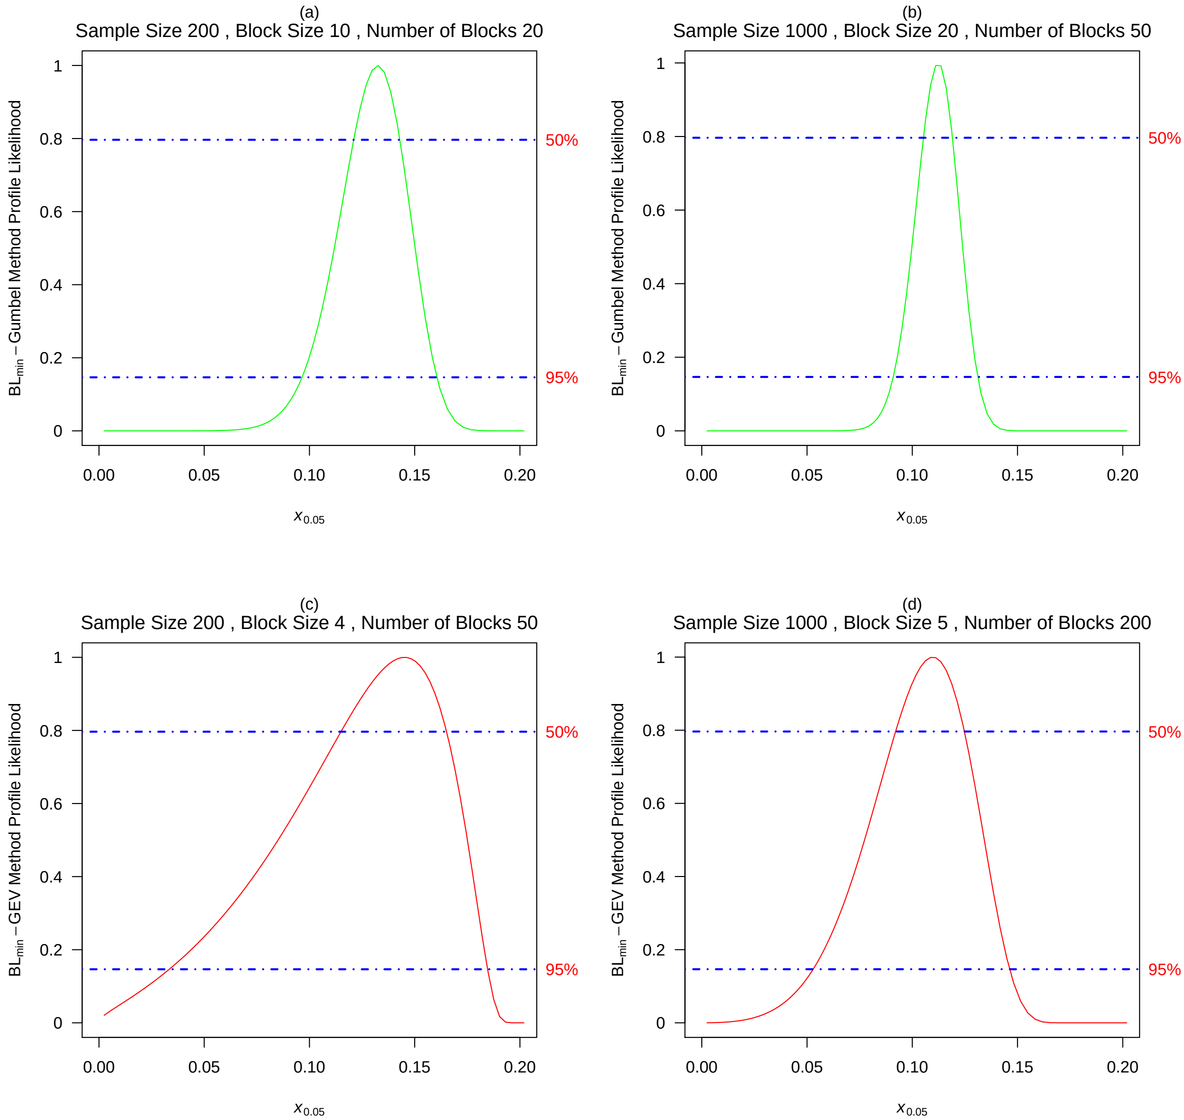

Figure 12: Relative likelihood profile plots for the quantile $x_{0.05}$ under the $\mathrm{BL}_{\mathrm{min}}$-Gumbel and the $\mathrm{BL}_{\mathrm{min}}$-GEV methods for $n = 200$ and $n = 1{,}000$. The horizontal lines indicate corresponding approximate 50% and 95% confidence intervals.

probability paper under the different estimation methods when sample sizes are $n = 200$ (left) and $n = 1{,}000$ (right).

Recall that Figure 2 suggested that the parent distribution of wall thicknesses could be described by either a Weibull or a lognormal distribution, but that the Weibull distribution fits better. A similar plot for the population of 5,649 thickness values (not shown here) confirmed that the Weibull provides a better description than the lognormal distribution. Among many distributions we tried, the Weibull distribution fits the 5,649 population thickness values best, especially in the lower tail. We then use the Weibull distribution to describe population parent distribution. Using all of the $M = 5{,}649$ thickness values to estimate the parameters in the Weibull parent distribution, the ML estimate for the 0.05 quantile of the minimum Weibull distribution is 0.115. When comparing estimates for different models, we use this quantity as the "true quantile" being estimated. Also the actual minimum thickness of the overall 5,649 features in the population was 0.146 inches, indicated in the plots by a vertical line.

In the top two plots in Figure 13, comparing the distribution estimates for the Weibull and lognormal parent distributions, we see the strong divergence from the truth when extrapolating toward small probabilities. Also, the estimate of the distribution of the minimum based on the lognormal assumption has serious upward bias. The estimate based on the Weibull distribution assumption is, as expected, close to the population quantile.

In the middle row of plots in Figure 13 we see that the $\mathrm{BL}_{\mathrm{min}}$-Gumbel method provides estimates of the distribution of the minimum that are similar to those of the minimum Weibull method, but with less precision due to the smaller effective sample size (number of blocks). Although this method would be expected to provide more robustness, as we saw in Section 5.4, with blocks of size 10 (used for the $n = 200$ sample), the degree of robustness would be limited if the parent distribution is misspecified as a lognormal.

In the bottom row of Figure 13 we see that the $\mathrm{BL}_{\mathrm{min}}$-GEV method also provides estimates of the distribution of the minimum that are similar to those of the minimum Weibull method, but now with much less precision, again due to the smaller effective sample size. Nevertheless, as we saw in Section 5.4, this method will not be affected by the model-specification bias that could affect the other estimation methods.

In all plots in Figure 13, the "true quantile" falls within the likelihood-based approximate 95% confidence intervals for $\widehat{x}_{0.05}$ under the three different estimation methods.

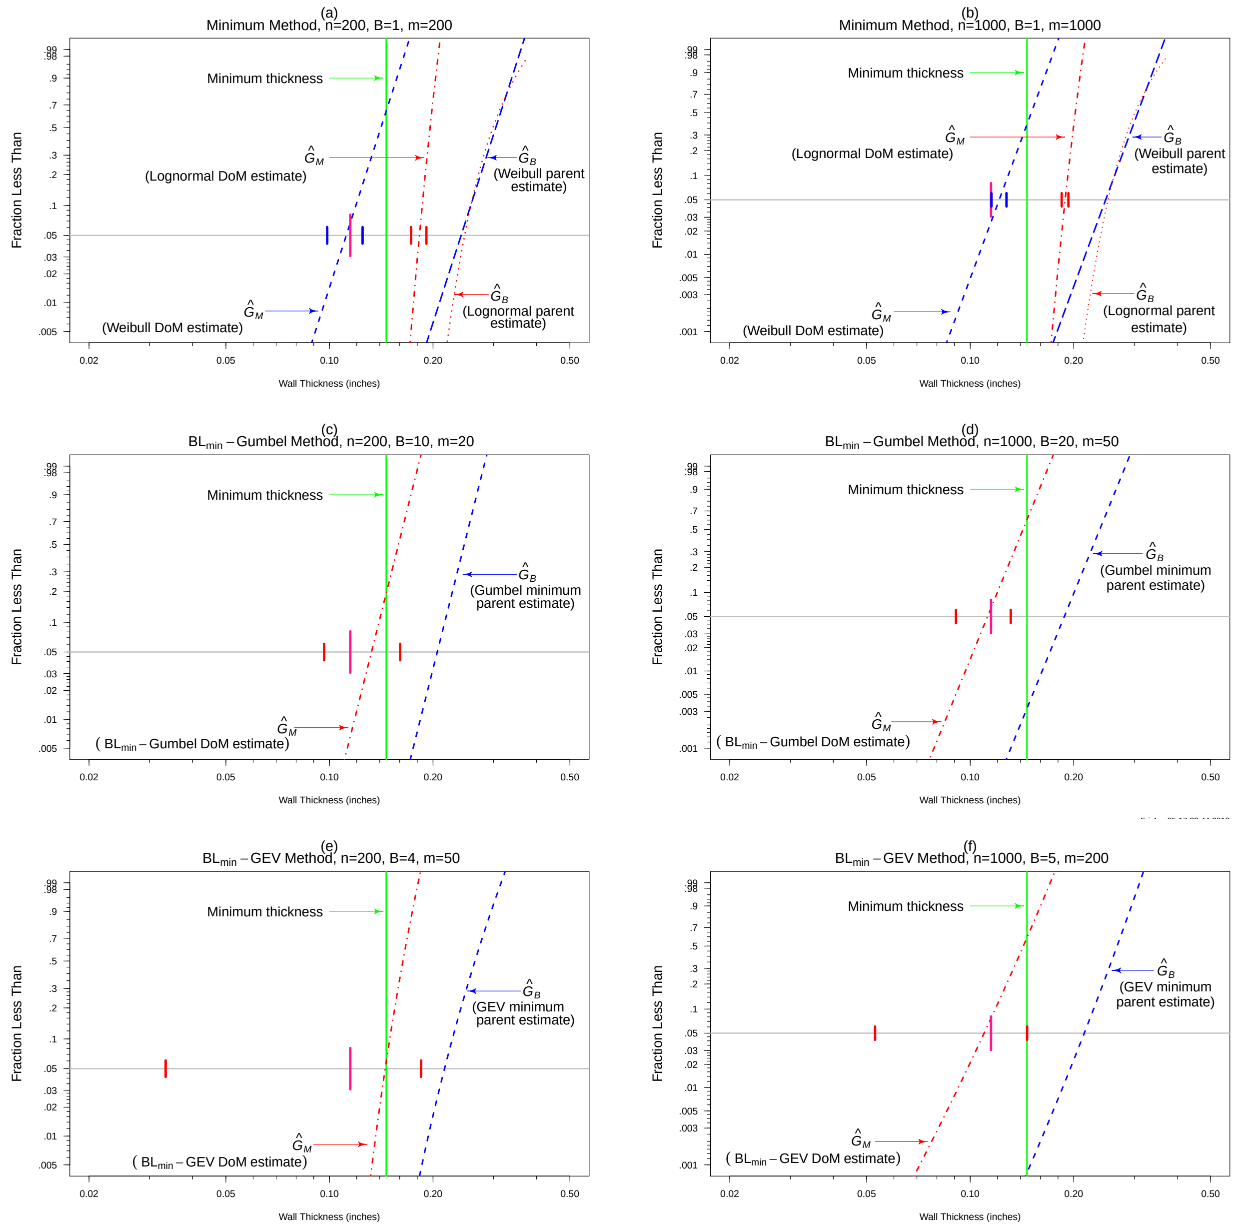

Figure 13: Weibull probability plots of estimates of the parent distribution and distribution of minimum (DoM) for the pipeline wall thickness inspection data under the three different estimation methods for $n = 200$ (left) and $n = 1{,}000$ (right). The shorter vertical tick marks on the $p = 0.05$ line indicate likelihood-based 95% confidence interval for the $x_{0.05}$ quantile of the minimum distributions. The longer vertical tick mark on the $p = 0.05$ line indicates "true quantile" $x_{0.05}$ based on the 5,649 population thickness values.

8 Conclusions, Recommendations, and Areas for Further Research

The observations from previous sections lead to the following conclusions:

• The direct distribution-of-minimum method provides the most precise quantile estimators when the parent distribution is correctly specified. Of course, in practice, we cannot expect, in many applications, to know the parent distribution exactly and we have seen that misspecification can lead to seriously biased answers.

• Sensitivity analysis showed that the performances of the incorrectly specified minimum distributions differ and depend highly on the lower tail behavior of the underlying parent distribution. For the normal parent distribution, the minimum distribution based on the SEV parent distribution performs poorly while the minimum distribution based on the LEV parent distribution provides fairly good quantile estimators. For the LEV parent distribution, although neither the SEV nor the normal minimum distributions provides good quantile estimators, the normal minimum distribution performs much better than the minimum distribution based on the SEV parent distribution. For the SEV parent distribution, both the normal and the LEV minimum distributions perform poorly. Compared with the performance of the LEV minimum distribution, the normal minimum distribution performs somewhat better. Generally, when extrapolating into the lower tail of the distribution, with the minimum distribution method using the SEV distribution assumption is conservative, relative to the normal and LEV distributions. This, of course, is not surprising given the lower tail behavior of these three distributions.

• In the block minima method, the choice of block size can be viewed as a trade-off between variance and bias and the trade-off is stronger for the $\mathrm{BL}_{\mathrm{min}}$-Gumbel method than it is for the $\mathrm{BL}_{\mathrm{min}}$-GEV method. For example, with a normal distribution parent, a large proportion of the MSE is contributed by the squared bias term, especially when the block sizes are small and the $\mathrm{BL}_{\mathrm{min}}$-Gumbel method is used. With a large number of blocks, the variance of the quantile estimator for the limiting distribution is relatively small. With a fixed amount of data, however, increasing the number of blocks will result in smaller blocks which will increase bias in the $\mathrm{BL}_{\mathrm{min}}$-Gumbel method (unless the parent is SEV), because the asymptotic extreme value theory assumes minima from large blocks.

• If there is a sufficiently large number of blocks and the blocks are of sufficient size, the $\mathrm{BL_{min}}$-GEV method provides inferences on the distribution of the minimum that are robust without need to specify the particular form of the parent distribution. If, however, the number of blocks is too small (say less than 30), ML estimation of the GEV parameter can fail to converge properly, no matter what the initial values are in the numerical optimization algorithm (Coles and Dixon 1999). Because of the unboundedness of the usual (product of densities) GEV likelihood, we used the "correct" (probability based) likelihood (as described, for example, by Barnard 1967, and Giesbrecht and Kempthorne 1976) to estimate the GEV parameters. As seen in Figure 9, even with this approach, large positive shape parameter estimates can arise and these result in the extremely small quantile estimates and in the $\mathrm{BL_{min}}$-GEV method, the degree of this behavior depends strongly on the shape of the parent distribution.

• In the $\mathrm{BL_{min}}$-Gumbel method, if the parent distribution is close to the lognormal or the Fréchet distribution (normal or LEV on the log scale), a large block size is needed to provide quantile estimates with small MSEs. If, however, the lower tail of the parent distribution is close to that of the Weibull distribution (SEV on the log scale), the $\mathrm{BL_{min}}$-Gumbel method with a small block size performs well, even if the block sizes are small.

• Other simulation results (details not shown here) show that the $\mathrm{BL_{min}}$-GEV method will also, in general, have some bias due to the finite block size but the bias is small relative to that in the $\mathrm{BL_{min}}$-Gumbel method and unless the number of blocks is large, the MSE tends to be dominated by variance.

• In the $\mathrm{BL_{min}}$-GEV method, if the parent distribution is close to the lognormal or the Fréchet distribution (normal or LEV on the log scale), using a large number of blocks (even with a small block size) provides good quantile estimates. If the parent distribution is close to the Weibull distribution (SEV on the log scale), the number of blocks needed to give reasonable precision is larger.

• For a given sample size and quantile of the minimum distribution, the ratio M/n affects the effective amount of extrapolation. As the ratio M/n increases (implying more extreme extrapolation into the distribution tail), bias and variance in the estimators of the quantile of interest will increase for both the $\mathrm{BL_{min}}$-GEV method and the $\mathrm{BL_{min}}$-Gumbel method.

Based on these conclusions, we have the following recommendations:

• The performance of ML estimators for quantiles based on the $\mathrm{BL_{min}}$-GEV method can be poor unless there is a large number of blocks (say greater than 50). One could use alternative point estimation methods. For example, the probability weighted moments method (PWM) described by Hosking, Wallis and Wood (1985) have been shown to have small sample superiority. It is not clear, however, that such alternative methods offer improvement when it is necessary to find confidence intervals for the quantile of interest (as is generally the case).

• Choosing an appropriate block size is essential for the successful use of the block minima method. Because the tail behaviors of different parent distributions could result in different shape parameter estimates in the GEV distribution, in order to make a decision between the $\mathrm{BL_{min}}$-GEV method and the $\mathrm{BL_{min}}$-Gumbel method, one should consider the degree of confidence that one has in the knowledge (perhaps due to physical knowledge or sampling considerations) of the parent distribution and the domain of attraction. After choosing a method, such knowledge is also important for choosing a block size.

• If the lower tail of the parent distribution can be appropriately described by the lognormal or the Fréchet distribution (normal or LEV on the log scale), one should choose a small block size with an adequately large number of the blocks when using the $\mathrm{BL_{min}}$-GEV method. The $\mathrm{BL_{min}}$-Gumbel method, however, requires a relatively larger block size to produce precise quantile estimators.

• If the parent distribution is closer to the Weibull distribution (SEV on the log scale), especially in the lower tail, the $\mathrm{BL_{min}}$-Gumbel method with a small block size for the quantile estimation is recommended.

• With a large sample size (e.g., n = 1,000), the $\mathrm{BL_{min}}$-GEV method provides an attractive method because of the robustness that it provides. For smaller sample sizes, the number of blocks used may not be large enough to provide a reasonable amount of precision. The $\mathrm{BL_{min}}$-Gumbel method is recommended for a moderate sample size (e.g., $n = 200$), but distribution-specification bias may be large if block-size distributions are not large enough. When the sample size is small (e.g., n = 20), one should use the traditional statistical method to model the minimum, recognizing that serious bias could be an issue if the parent distribution is seriously misspecified.

Some areas for future research are:

• Our study focuses on the evaluation of the point estimates. In statistical inference, the accuracy of the quantile estimators using different estimation methods is another main issue. It would be of interest to compare confidence intervals of the quantile estimates in terms of the coverage probability.

• Methods that relax the restriction of equal block size in the block minima method could provide flexibility in modeling the minimum thickness of the pipeline and choosing the appropriate block sizes in the block minima extreme value theory method.

• The peaks over threshold (POT) method is an alternative to the block minima. A study similar to this one could be conducted to investigate threshold choice for the POT method.

• Bayesian methods, combining the pipeline wall thickness measurements with the prior information, especially, on the GEV shape parameter, have the potential to provide more precise quantile estimate for the minimum GEV distribution when legitimate prior information is available.

• Method for analyzing pipeline data taken over time to estimate corrosion rates and predict minimum at future points in time could be developed.

9 Acknowledgments

We would like to thank Luis Escobar, Ulrike Genschel, Yili Hong and Alyson Wilson for providing helpful comments on an earlier version of this paper.

References

[1] Barlow, R. E. and Proschan, F. (1975). *Statistical Theory of Reliability and Life Testing: Probability Models*. New York: Holt, Rinehart and Winston.

[2] Barnard, G. A. (1967). The use of the likelihood function in statistical practice. In *Fifth Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pages 27-40. University of California Press.

[3] Castillo, E., Hadi, A. S., Balakrishnan, N., and Sarabia, J. M. (2005). *Extreme Value and Related Models with Applications in Engineering and Science*. New York: John Wiley and Sons.

[4] Coles, S. G. (2001). *An Introduction to Statistical Modeling of Extreme Values*. New York: Springer.

[5] Coles, S. G. and Dixon, M. (1999). Likelihood-based inference for extreme value models. *Extremes*, 2:5-23.

[6] Cox, D. R., Isham, V. S., and Northrop, P. J. (2002). Floods: some probabilistic and statistical approaches. *Philosophical Transactions of the Royal Society of London, Series A*, 360:1389-1408.

[7] Engeland, K., Hisdal, H., and Frigessi, A. (2004). Practical extreme value modelling of hydrological floods and droughts: A case study. *Extremes*, 7:5-30.

[8] Fisher, R. and Tippett, L. (1928). On the estimation of the frequency distributions of the largest or smallest member of a sample. *Proceedings of the Cambridge Philosophical Society.*, 24:180-190.

[9] Fougères, A. L., Holm, S., and Rootzén, H. (2006). Pitting corrosion: Comparison of treatments with extreme-value-distributed responses. *Technometrics*, 48:262-272.

[10] Giesbrecht, F. and Kempthorne, O. (1976). Maximum likelihood estimation in the three-parameter lognormal distribution. *Journal of the Royal Statistical Society B*, 38:257-264.

[11] Glegola, M. (2007). Extreme value analysis of corrosion data. Master's thesis, Delft University of Technology, Delft, Netherlands. `http://risk2.ewi.tudelft.nl/research-and-publications/doc_view/158-mglegolathesispdf`.

[12] Gumbel, E. J. (1958). *Statistics of Extremes*. New York: Columbia University Press.

[13] Hosking, J. R. M., Wallis, J. R., and Wood, E. F. (1985). Estimation of the generalized extreme-value distribution by the method of probability-weighted moments. *Technometrics*, 27:251-261.

[14] Kowaka, M., Tsuge, H., Akashi, M., Masamura, K., and Ishimoto, H. (1994). *Introduction to Life Prediction of Industrial Plant Materials : Application of the Extreme Value Statistical Method for Corrosion Analysis.* New York : Allerton Press.

[15] Laycock, P. J., Cottis, R. A., and Scarf, P. A. (1990). Extrapolation of extreme pit depths in space and time. *Journal of The Electrochemical Society*, 137: 64-69.

[16] Laycock, P. J. and Scarf, P. A. (1993). Exceedances, extremes, extrapolation and order statistics for pits, pitting and other localized corrosion phenomena. *Corrosion Science*, 35: 135-145.

[17] Lawless, J. F. (2002). *Statistical Models and Methods for Lifetime Data.* New York: John Wiley and Sons.

[18] Meeker, W. Q. and Escobar, L. A. (1998). *Statistical Methods for Reliability Data.* New York: John Wiley and Sons.

[19] Scarf, P. A. and Laycock, P. J. (1994). Applications of extreme value theory in corrosion engineering. *Journal of Research of the National Institute of Standards and Technology*, 99:313-320.

[20] Scarf, P. A. and Laycock, P. J. (1996). Estimation of extremes in corrosion engineering. *Journal of Applied Statistics*, 23:621-643.

[21] Shibata, T. (1994). Application of extreme value statistics to corrosion. *Journal of Research of the National Institute of Standards and Technology*, 99:327-336.

[22] Vardeman, S. and Lee, C. (2005). Likelihood-based statistical estimation from quantized data. *IEEE Transactions on Instrumentation and Measurement*, 54:409-414.