arXiv:1803.10766v3 [stat.ME] 22 Jul 2019

Repeated out of Sample Fusion in the Estimation of Small Tail Probabilities

Benjamin Kedem, Lemeng Pan, Paul Smith, and Chen Wang

Department of Mathematics
University of Maryland, College Park

May 2019

Abstract

Often, it is required to estimate the probability that a quantity such as toxicity level, plutonium, temperature, rainfall, damage, wind speed, wave size, earthquake magnitude, risk, etc., exceeds an unsafe high threshold. The probability in question is then very small. To estimate such a probability, information is needed about large values of the quantity of interest. However, in many cases, the data only contain values below or even far below the designated threshold, let alone exceedingly large values. It is shown that by repeated fusion of the data with externally generated random data, more information about small tail probabilities is obtained with the aid of certain new statistical functions. This provides relatively short, yet reliable interval estimates based on moderately large samples. A comparison of the approach with a method from extreme values theory (peaks over threshold, or POT), using both artificial and real data, points to the merit of repeated out of sample fusion.

Keywords: Density ratio model, semiparametric, coverage, iterative, peaks-over-threshold, B-curve.

MSC 2000: Primary 62F40; Secondary 62F25

1 Introduction

This paper addresses the following basic problem. Consider a moderately large random sample $\boldsymbol{X}_0 = (X_1, ..., X_{n_0})$ where all the observations are much smaller than a high threshold T, that is $\max(\boldsymbol{X}_0) << T$. Based on the sample we wish to estimate the probability p of exceeding T without knowing the underlying distribution. However, as is, the sample may not contain sufficient amount of information to tackle the problem. To gain more information, the problem is approached by combining or fusing the sample repeatedly with externally generated computer data.

The problem is inspired by real world situations where the measurements X_i are all below T. As an example, consider a sample of rogue ocean waves none of which exceeds $T = 150$ feet in height, and yet we wish to estimate the small chance of exceeding T from a moderately large sample $\boldsymbol{X}_0 = (X_1, ..., X_{n_0})$, referred to as a *reference* sample. Similar problems pertain to insurance claims, food safety, and environmental risks such as radiation levels.

How is this done? This article advances a statistical notion where a small tail probability is identified with a *point* on a certain monotone curve called *B-curve*, obtained by repeated fusion of real and artificial data. The point on the curve is approached by an iterative procedure, against the backdrop of numerous fusions of real and computer generated data, leading to a certain random "fixed point" for a lack of a better term.

Let $\boldsymbol{X}_i$ denote the ith computer generated sample of size $n_1 = n_0$. Then the *fused* or *combined* samples are the *augmentations*

$$(\boldsymbol{X}_0, \boldsymbol{X}_1), (\boldsymbol{X}_0, \boldsymbol{X}_2), (\boldsymbol{X}_0, \boldsymbol{X}_3)... \tag{1}$$

where $\boldsymbol{X}_0$ is a real reference sample and the $\boldsymbol{X}_i$ are different independent computer generated samples. The number of fusions can be as large as we wish. For example 10,000 or 100,000 or 1,000,000 or more fusions.

In that way, by repeated fusion, we can extract information about tail probabilities possibly not available in the original reference sample by itself. As we shall see, in many cases this brings about surprisingly precise estimates for small tail probabilities, using moderately large samples (e.g. 100 or 200 data points), as described and illustrated in Section 4.

Fusing a given sample repeatedly with computer generated data is referred to as *repeated out of sample fusion (ROSF)* (Kedem et al. 2016). Unlike the bootstrap, additional information is sought repeatedly from outside the sample. Related ideas concerning a single fusion are studied in

Fithian and Wager (2015), Fokianos and Qin (2008), Katzoff et al. (2014), and Zhou (2013).

1.1 An Iterative Procedure and its Estimates

As said, the large number of fusions results in a *B-curve* defined in Section 2. The B-curve is monotonically increasing and it contains a *point* whose ordinate is very close to p with a high probability. As the number of fusions increases the ordinate of that point essentially coincides with p. The goal is to "capture" that point by an iterative algorithm.

The consequential estimates as well as interval estimates of p are quite precise. A comparison with peaks-over-threshold (POT) from extreme value theory (Beirlant et al. 2004, Ferreira and DeHaan 2015) indicates that ROSF can bring about a substantial gain in reliability as well as in precision across a fairly wide range of tail behavior, given moderately large samples $\boldsymbol{X}_0$.

The question then is how to tie or connect the real data and the generated random data to obtain useful reliable estimates for small tail probabilities. Connecting or fusing the real and artificial data can be approached by means of their respective probability distributions under the so called *density ratio model* framework, discussed briefly in Section 3 and in the Appendix.

Thus, the paper describes ROSF and a related iterative method (IM) in the estimation of small tail probabilities, against the backdrop of the density ratio model, by "capturing" a point on the B-curve.

1.2 A Typical "Down-Up" Example

We illuminate our method upfront by a typical example, postponing technical details to later sections. Let $\boldsymbol{X}_0$ be a lognormal random sample LN(1,1), of size $n_0 = 100$. The largest observed data point was $\max(\boldsymbol{X}_0) = 32.36495$, while $T = 59.75377$. Hence, $\max(\boldsymbol{X}_0) = 32.36495 << T$. We wish to estimate the probability of exceeding T. That is, we wish to estimate the tail probability $p = P(X > T)$, which in the present case is $p = 0.001$.

As in (1), we fused $\boldsymbol{X}_0$ repeatedly with 10,000 computer generated Unif(0,80) samples of size $n_1 = n_0 = 100$. The iterative method produces "D=Down" and "U=Up" sequences. When the "Down" transitions to "Up" we know we are very close to the true p. In the present example we have:

D 0.001199466, D 0.001099466, D 0.000999465, D 0.000999465,
D 0.000999465, D 0.000999465, U 0.000999465, U 0.000999465.

Here the transition from Down to Up occurs at 0.000999465, which can be taken as point estimate for $p = 0.001$. In this example the error is quite small.

This "Down-Up" phenomenon, giving surprisingly precise estimates, has been observed numerous times across many different tail types as we shall demonstrate in Tables 3 to 14 in Section 4.2, where we also return to the present example providing more details.

2 ROSF and the B-Curve

We are in pursuit of a small tail probability p. It is shown how to construct a curve which contains with a high probability a point whose ordinate is p.

Suppose $\boldsymbol{X}_0 = (X_1, ..., X_{n_0})$ is a reference sample from some reference distribution g, and that we wish to estimate a small tail probability p of that distribution. The $X_i \sim g$ could represent quantities such as earthquake magnitude, radioactive contamination, claim amount, financial return, poverty level, wealth, temperature, and so on, and the interest is in the tail probability $p = P(X > T)$ for some relatively high threshold T.

Combining the reference sample $\boldsymbol{X}_0$ with a computer-generated sample $\boldsymbol{X}_1$ gives the fused sample $(\boldsymbol{X}_0, \boldsymbol{X}_1)$. Then $\boldsymbol{X}_0$ can be fused again with another independent computer generated sample $\boldsymbol{X}_2$ and we get another fused sample $(\boldsymbol{X}_0, \boldsymbol{X}_2)$, and so on. All these computer-generated samples $\boldsymbol{X}_1$, $\boldsymbol{X}_2$,... are independent and are generated in an identical manner, and all have the same size n_1. We refer to these computer-generated samples as *fusion samples*. Observe that the fused or combined samples $(\boldsymbol{X}_0, \boldsymbol{X}_1)$, $(\boldsymbol{X}_0, \boldsymbol{X}_2)$,... all have size $n_0 + n_1$.

Here is how B-curves are constructed. We fuse the given reference sample $\boldsymbol{X}_0$ with a computer-generated fusion sample $\boldsymbol{X}_1$ from g_1 and get in a certain way, described in the next section, a confidence interval for the small tail probability p. Let B_1 denote the upper bound of that interval. We fuse the given reference sample $\boldsymbol{X}_0$ again with another artificial fusion sample $\boldsymbol{X}_2$ from g_1 and get in the same manner another upper bound B_2 for p. This process is repeated many times to produce a long sequence of confidence intervals and hence a long sequence of upper bounds B_i, $i = 1, 2, ...$ Conditional on $\boldsymbol{X}_0$, the sequence of upper bounds $B_1, B_2, ...$ is then an independent and identically distributed sequence of random variables from some distribution F_B. It is assumed that

$$P(B_1 > p) = 1 - F_B(p) > 0. \tag{2}$$

Let $B_{(1)}, B_{(2)}, ..., B_{(N)}$ be a sequence of order statistics from smallest to largest. Then, as $N \to \infty$, $B_{(1)}$ decreases and $B_{(N)}$ increases. Hence, as the number of fusions increases the plot consisting of the pairs

$$(1, B_{(1)}), (2, B_{(2)}), ..., (N, B_{(N)}) \tag{3}$$

contains a point whose ordinate is p with probability approaching 1. It follows that as $N \to \infty$, there is a $B_{(j)}$ which essentially coincides with p. The plot of points consisting of the pairs $(j, B_{(j)})$ in (3) is referred to as the *B-curve*.

Typical B-curves corresponding to the tail probability $p = P(X > T) = 0.001$ for various reference samples $\boldsymbol{X}_0$ from the indicated distributions or data are shown in Figure 1. Notice that to get $p = 0.001$, in each case the threshold T must change accordingly, and that in each plot there is a $B_{(j)}$ nearest or closest to $p = 0.001$. The curves were obtained from 10,000 fusions of $\boldsymbol{X}_0$ with uniform samples with support exceeding T. Clearly, in all cases $B_{(1)} < p < B_{(10,000)}$.

Figure 2 shows B-curves from the f(2,7) distribution for various $\max(\boldsymbol{X}_0)$ where $T = 21.689$. The point "•" moves to the left as $\max(\boldsymbol{X}_0)$ increases relative to $T = 21.689$. We can see that in each plot there is a $B_{(j)}$ nearest or closest to $p = 0.001$. As before, the curves were obtained from 10,000 fusions of $\boldsymbol{X}_0$ with uniform samples with support exceeding T, and in all cases $B_{(1)} < p < B_{(10,000)}$.

A key fact of the present approach is that since the fusions can be repeated indefinitely, we can approximate the distribution of the B upper bounds arbitrarily closely.

Let $\hat{F}_B$ be the empirical distribution obtained from the sequence of upper bounds $B_1, B_2, ..., B_N$. Then from the Glivenko-Cantelli Theorem, $\hat{F}_B$ converges to F_B almost surely uniformly as N increases. Since the number of fusions can be as large as we wish, *our key idea*, F_B is known for all practical purposes.

3 Getting Upper Bounds by Data Fusion

This section describes a particular way of generating upper bounds for tail probabilities p by data fusion of the real $\boldsymbol{X}_0$ and additional computer-generated data ("augmented reality" as it were) under the density ratio model defined in (5) below. The upper bounds are needed in order to generate B-curves.

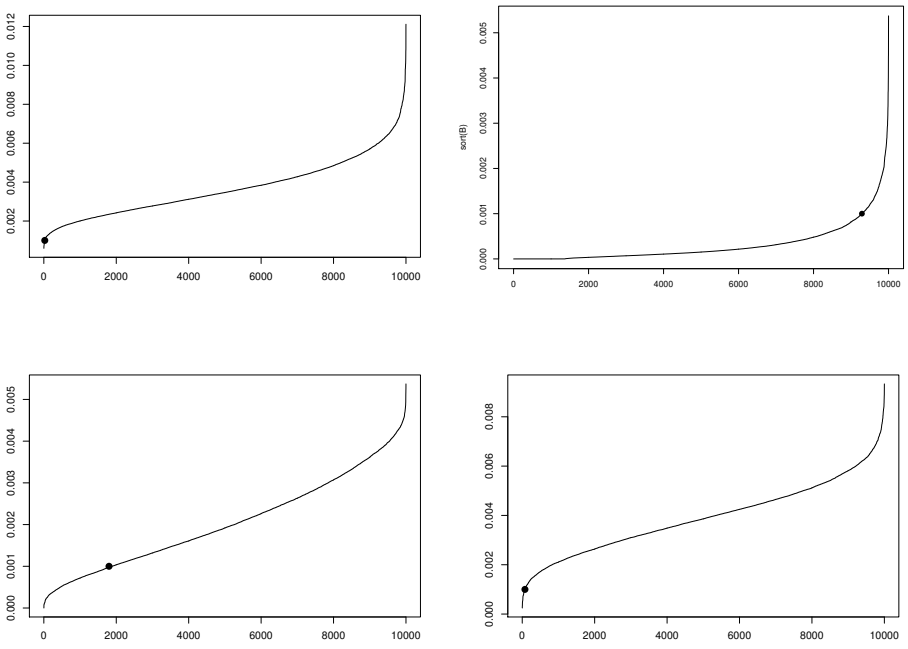

Figure 1: Typical B-Curves from $B_{(1)}, ..., B_{(10{,}000)}$ containing a point corresponding to $p = 0.001$. Clockwise from top left: Gamma(1,0.01), LN(1,1), Lead exposure, Mercury. T=690.7755, 59.7538, 25.00, 22.41, respectively, $n_0 = n_1 = 100$. In all cases the fusion samples are uniform with support exceeding T.

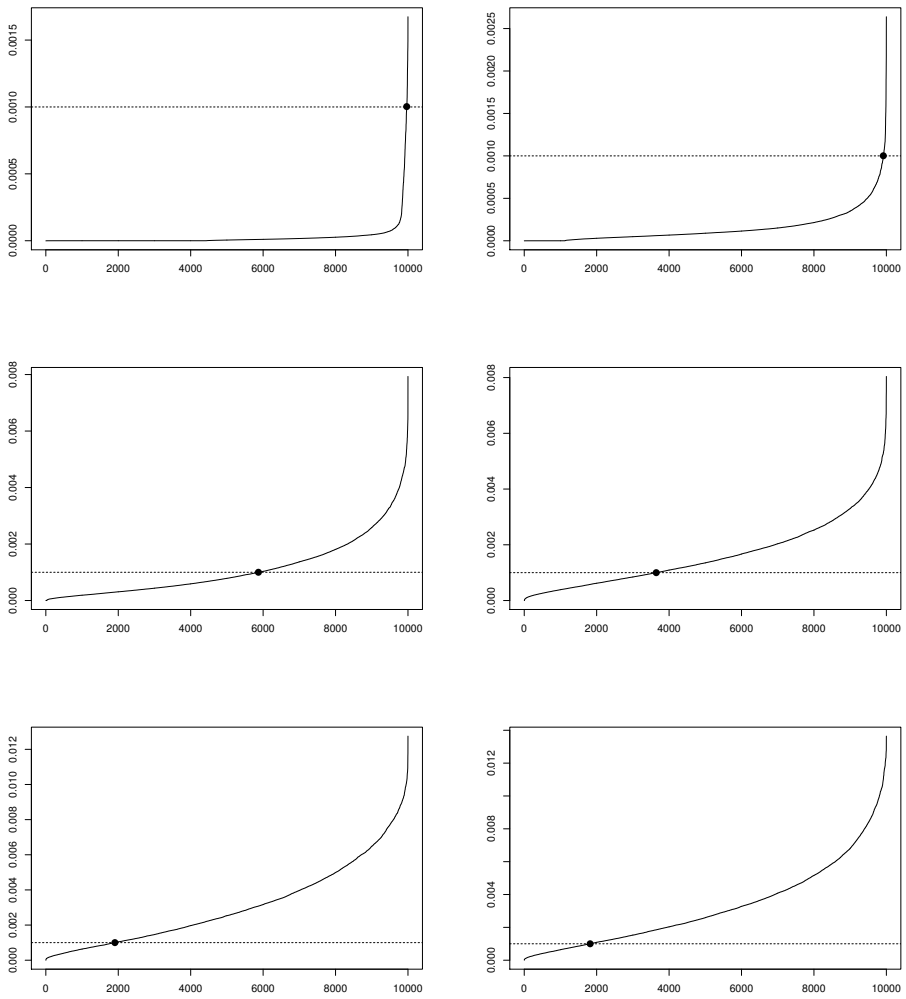

Figure 2: B-Curves, 10,000 B's, from f(2,7). $n_0 = n_1 = 100$, $p = 0.001$, $T = 21.689$. $\max(\boldsymbol{X}_0)$ values: Top left 4.111466, top right 6.86562. Middle left 8.631132, middle right 10.18021. Bottom left 17.26258, bottom right 19.95937. The point "•" moves to the left as $\max(\boldsymbol{X}_0)$ increases relative to $T = 21.689$. The fusion samples are uniform with support covering T.

In general, by "fusion" or "data fusion" we mean the combined data from $m+1$ sources where each source is governed by a probability distribution. In the spirit of augmented reality, computer algorithms which generate random data are perfectly legitimate data sources. Using the combined data, semiparametric statistical inference can be ensued under the density ratio model assumption (Kedem, et al. 2017, Lu 2007, Qin and Zhang 1997).

Assume that the reference random sample $\boldsymbol{X}_0$ of size n_0 follows an unknown reference distribution with probability density g, and let G be the corresponding cumulative distribution function (cdf).

Let

$$\boldsymbol{X}_1, ..., \boldsymbol{X}_m,$$

be additional computer-generated random samples where $\boldsymbol{X}_j \sim g_j, G_j$, with size n_j, $j = 1, ..., m$. As in the Appendix, for now $m \geq 1$ but later we specialize to $m = 1$ only as in (1). The augmentation of $m+1$ samples

$$\boldsymbol{t} = (t_1, \ldots, t_n) = (\boldsymbol{X}_0, \boldsymbol{X}_1, \ldots, \boldsymbol{X}_m), \tag{4}$$

of size $n_0 + n_1 + \cdots + n_m$ gives the fused data. The density ratio model stipulates that

$$\frac{g_j(x)}{g(x)} = \exp(\alpha_j + \boldsymbol{\beta}_j' \boldsymbol{h}(x)), \quad j = 1, \ldots, m, \tag{5}$$

where $\boldsymbol{\beta}_j$ is an $r \times 1$ parameter vector, α_j is a scalar parameter, and $\boldsymbol{h}(x)$ is an $r \times 1$ vector valued distortion or tilt function. Clearly, to generate the $\boldsymbol{X}_j$ samples we must know the corresponding g_j. However, beyond the generating process, we do not make use of this knowledge. Thus, by our estimation procedure, none of the probability densities $g, g_1, ..., g_m$ and the corresponding G_j's, and none of the parameters α's and $\boldsymbol{\beta}$'s are assumed known, but, strictly speaking, the so called tilt function $\boldsymbol{h}$ must be a known function. However, in the present application the requirement of a known $\boldsymbol{h}$ is apparently mitigated as accentuated by assumption (2) above, which may hold for misspecified $\boldsymbol{h}$, and by our examples with many different tail types.

Since all the probability distributions are connected by the density ratio model (5), each distribution pair g_j, G_j is estimated from the entire fused data $\boldsymbol{t}$ and not just from $\boldsymbol{X}_j$ only. The same holds for the reference pair g, G. Thus, for example, the reference G is estimated from the entire fused data $\boldsymbol{t}$ with $n_0 + n_1 + \cdots + n_m$ observations and not only from the reference sample $\boldsymbol{X}_0$ with n_0 observations.

Under the density ratio model (5), the maximum likelihood estimate of $G(x)$ based on the fused data $\boldsymbol{t}$ is given in (12) in Section A.1 in the Appendix, along with its asymptotic distribution described in Theorem A.1. From the theorem we obtain confidence intervals for $p = 1 - G(T)$ for any threshold T using (15). In addition, from (12) we get the point estimates for p as well. It is different than the one obtained by the "Down-Up" idea, and is not used in the paper, as in many cases it underestimates p.

Our data analysis in Section 4 regarding many different tail types, and additional examples in Kedem et al. (2017), indicate that for the implementation of ROSF, the density ratio model need not hold precisely, and that the "gamma tilt" $\boldsymbol{h}(x) = (x, \log x)$ is a sensible choice for skewed data. In that case, many of the B_i obtained from (15) will be greater than p as their number increases, but some will not. Hence, the *ordered* $B_{(i)}$ engulf or surround p with probability approaching one as the number of fusions increases. That is, as the number of fusions increases, the set of pairs $(j, B_{(j)})$ engulfs the desired point on the B-curve with probability approaching one. This is illustrated in Figures 1 and 2 with 10,000 fusions.

This, in general, cannot be said about the ordered $\hat{p}$'s from (12) unless the number of fusions is exceedingly large. Hence, we shall not use $\hat{p} = 1 - \hat{G}(T)$ from (12). Instead we estimate p from the iterative process described earlier and in Section 4.

In this paper $m = 1$ only, and the fusion samples are uniform random samples supported over a wide range which covers T. The reason for uniform samples is that when the density ratio model holds for some g and g_1, then it also holds approximately by taking g_1 as a uniform density supported over a sufficiently wide range.

To summarize, numerous examples with skewed data suggest that the confidence intervals (15) are still useful in conjunction with $\boldsymbol{h}(x) = (x, \log x)$ even when the density ratio model does not hold in a strict sense. In that case, the reference sample $\boldsymbol{X}_0$ is fused repeatedly with identically distributed independent random uniform samples $\boldsymbol{X}_1$, $\boldsymbol{X}_2$,..., as in (1), where the upper limit of the uniform support exceeds T. Repeated fusion gives upper bounds $B_1, B_2, ...$ for $p = P(X > T)$ using (15). Conditional on $\boldsymbol{X}_0$, the upper bounds B_i are independent and identically distributed random variables from some distribution F_B. The B-curves are constructed from a large number of pairs $(j, B_{(j)})$, where the $B_{(j)}$ are the ordered upper bounds.

4 Capturing a Point on the B-Curve

Due to a large number of fusions n, F_B is known for all practical purposes and with probability close to 1

$$B_{(1)} < p < B_{(n)}. \tag{6}$$

In general, even for $n = 1,000$, $B_{(1,000)}$ is much larger than the true p and $B_{(1)}$ is very close to 0. The goal is to find $B_{(j)}$ close to p.

It follows, by the monotonicity of the B-curve and (6), that as j *decreases* (for example from $n = 10,000$), the $B_{(j)}$ approach p from above so that there is a $B_{(j)}$ very close to p. Likewise, the $B_{(j)}$ can approach p from below. Thus, the B-curve establishes a relationship between j and p.

Another relationship between j and p is obtained from a basic fact about order statistics where it is known that

$$P(B_{(j)} > p) = \sum_{k=0}^{j-1} \binom{n}{k} [F_B(p)]^k [1 - F_B(p)]^{n-k}. \tag{7}$$

Suppose now that the probabilities

$$P(B_{(j_1)} > p_{j_1}),\ P(B_{(j_2)} > p_{j_2}), ...$$

are sufficiently high probabilities, and that from the B-curve we get the close approximations

$$p_{j_1} \doteq B_{(j_2)},\ p_{j_2} \doteq B_{(j_3)}....$$

Then with a high probability we get a decreasing "down" sequence

$$B_{(j_1)} > B_{(j_2)} > B_{(j_3)} \cdots .$$

Replacing the "sufficiently high probabilities" by "sufficiently low probabilities", then a dual argument leads to an increasing "up" sequence

$$B_{(j_1')} < B_{(j_2')} < B_{(j_3')} \cdots .$$

Thus, when the probabilities are sufficiently high the $B_{(j_k)}$ decrease, and when the probabilities are sufficiently low the $B_{(j_k)}$ increase. In particular, this "Down-Up" phenomenon occurs in a neighborhood of the true p, where a *transition or shift* occurs from "down" to "up" or vice versa, resulting in a "capture" of p. Thus, allowing for high and low probabilities by bounding (7) by a sufficiently high probability, we have.

Proposition: *Assume that the samples size n_0 of $\boldsymbol{X}_0$ is large enough, and that the number of fusions n is sufficiently large so that $B_{(1)} < p < B_{(n)}$. Consider the smallest $p_j \in (0,1)$ which satisfy the inequality*

$$P(B_{(j)} > p_j) = \sum_{k=0}^{j-1} \binom{n}{k} [F_B(p_j)]^k [1 - F_B(p_j)]^{n-k} \leq 0.95, \tag{8}$$

where the p_j are evaluated along appropriate numerical increments. Then, (8) produces "down" and "up" sequences depending on the $B_{(j)}$ relative to p_j. In particular, in a neighborhood of the true tail probability p, with a high probability, there are "down" sequences which converge from above and "up" sequences which converge from below to points close to p.

This will be demonstrated copiously across different tail types using an approximation to (8). We note that (7) is a steep monotone decreasing step function so that if "$>$" is used instead of "$\leq$" in (8) then the solution of (8) is $p = 0$, and that replacing 0.95 by 0.99 in (8) gives similar results.

Iterating between these two monotone relationships, the B-curve and (8), is what was referred to earlier as the iterative method (IM). The iterative method provides our p estimates. The iterations could start with a sufficiently large j suggested by the B-curve, or, alternatively with a sufficiently small j, until the down and up sequences converge to the same or very close points. The average of these points, or an approximation thereof, is our point estimate from the iterative process and it is different than $\hat{p}$ obtained from (12) in the Appendix.

In general, starting with any j, convergence occurs by monotonicity and we keep getting the same point.

In symbols, with $B_{(j_k)}$'s from the B-curve, and $p_{(j_k)}$'s the smallest p's satisfying (8) with $j = j_k$, and $B_{(j_{k+1})}$ closest to $p_{(j_k)}$, $k = 1, 2, ...,$

$$B_{(j_1)} \to p_{(j_1)} \to B_{(j_2)} \to \cdots B_{(j_k)} \to p_{j_k} \to B_{(j_{k+1})} \to p_{j_k} \to B_{(j_{k+1})} \to p_{j_k} \cdots$$

so that p_{j_k} keeps giving the same $B_{(j_{k+1})}$ (and hence the same j_{k+1}) and vice versa. This can be expressed more succinctly as,

$$j_1 \to p_{(j_1)} \to j_2 \to p_{(j_2)} \to \cdots j_k \to p_{j_k} \to j_{k+1} \to p_{j_k} \to j_{k+1} \to p_{j_k} \cdots$$

As will be illustrated in Section 4.2, under some computational conditions this iterative process results in a contraction in a neighborhood of the true p.

4.1 Computational Considerations

Computationally, the iterative process depends on n and the increments of p at which (8) is evaluated. In practice, due to computational limitations of large binomial coefficients the iteration is done as follows. After F_B is obtained from a large number of fusions, say $n = 10,000$ fusions (which give 10,000 B's), then 1000 B's are sampled at random from the obtained $n = 10,000$ B's to get an approximate B-curve. Next, the binomial coefficients $\binom{n}{k}$ are replaced by $\binom{1000}{k}$. We then iterate between an approximate B-curve and approximate (8) with $n = 1000$ as in

$$\sum_{k=0}^{j-1} \binom{1000}{k} [F_B(p_j)]^k [1 - F_B(p_j)]^{n-k} \leq 0.95 \tag{9}$$

until a "Down-Up" convergence occurs, in which case an estimate for p is obtained as the Down-Up shift point. The iterative process is illustrated in the next section. This procedure can be repeated many times by sampling repeatedly many different sets of 1000 B's to obtain many point estimates $\hat{p}$ from which interval estimates can then be constructed, as well as variance estimates.

Running 10,000 fusions takes about 5 minutes in R which translates into about 8 hours for 1,000,000 fusions. In what follows the p-increments at which (9) is evaluated are 0.0001 when $p = 0.001$ and 0.000015 when $p = 0.0001$. In all cases the maxima (minima) of the approximate B-curves occurred at a point larger (smaller) than the true p, as in (6).

4.2 Illustrations of an Iterative Process

The Down-Up convergence results together with the number of iterations are summarized in Tables 1 to 10 for $p = 0.001$ and also for $p = 0.0001$. Due to insufficiently large data sets, the Down-Up convergence results for real data in Tables 11 – 14 do not deal with the smaller $p = 0.0001$. The cdf F_B was obtained from 10,000 B's (the result of 10,000 fusions), and each entry in the tables was obtained from a *different* sample of 1,000 B's sampled at random from 10,000 B's.

4.2.1 Some Typical Down-Up Sequences

It is instructive first to realize some typical Down-Up sequences. In the first example $\boldsymbol{X}_0$ is a LN(1,1) sample where $\max(\boldsymbol{X}_0) = 32.36495$. With $T = 59.75377$ the true tail probability to be estimated is $p = 0.001$, using

$n_0 = n_1 = 100$ and $\boldsymbol{h} = (x, \log x)$. The generated fusion samples $\boldsymbol{X}_1$ are from Unif(0,80), $80 > T$, and F_B was obtained from 10,000 fusions.

Typical convergent Down-Up sequences (j, p_j) are given next. Again, each sequence was derived from a different B-sample of size 1000 drawn from 10,000 B's. More examples are given in Kedem et al. (2018).

Down: $900 \to 0.001799466 \to 867 \to 0.001599466 \to 837 \to 0.001499466 \to 822 \to 0.001399466 \to 801 \to 0.001299466 \to 778 \to 0.001199466 \to 751 \to 0.001099466 \to 723 \to 0.001099466 \cdots$,

Down: $800 \to 0.001299466 \to 775 \to 0.001199466 \to 743 \to 0.001099466 \to 712 \to 0.0009994658 \to 680 \to 0.0009994658 \cdots$.

Up: $680 \to 0.0009994658 \to 694 \to 0.0009994658 \cdots$,

Up: $670 \to 0.0008994658 \to 675 \to 0.0009994658 \to 711 \to 0.0009994658 \cdots$.

In the next example $\boldsymbol{X}_0$ is a mercury sample (see Section 4.2.8), drawn from a large population, where $\max(\boldsymbol{X}_0) = 7.99$. With $T = 22.41$ the true tail probability to be estimated is $p = 0.001088797$, using $n_0 = n_1 = 100$ and $\boldsymbol{h} = (x, \log x)$. The generated fusion samples $\boldsymbol{X}_1$ are from Unif(0,50), $50 > T$, and F_B was obtained from 10,000 fusions.

Typical convergent Down-Up sequences (j, p_j) are:

Down: $600 \to 0.001299352 \to 563 \to 0.001199352 \to 526 \to 0.001099352 \to 502 \to 0.0009993515 \to 475 \to 0.0009993515 \cdots$,

Down: $550 \to 0.001199352 \to 542 \to 0.001099352 \to 509 \to 0.0009993515 \to 479 \to 0.0009993515 \cdots$,

Up: $490 \to 0.0009993515 \to 503 \to 0.0009993515 \cdots$,

Up: $470 \to 0.0008993515 \to 476 \to 0.0009993515 \to 505 \to 0.0009993515 \cdots$.

We note that the number of Down-Up iterations decreases in a neighborhood of the true p. As seen from Tables 1 to 14 below, in many cases 1 or 2 iterations in a neighborhood of p suffice. This reduction can serve as a telltale sign that convergence took place.

We further note that the Gamma cases in Tables 1 and 2 are nearly specified whereas this cannot be said about the cases in Tables 3 to 14. However, the tables portray a very similar picture for both real and simulated data, for (nearly) specified or misspecified cases, giving precision on the order of 10^{-5} or better for $p = 0.001$ and order of 10^{-6} for $p = 0.0001$, where $n_0 = n_1 = 100$. The results in the tables were obtained from 10,000 out of sample fusions, and in all cases it has been observed that $p \in (B_{(1)}, B_{(10,000)})$ as also illustrated in Figures 1 and 2. *Notably, as seen from Tables 1 to 14, the Down-Up shift point is close to p.*

4.2.2 Gamma(1,0.05)

Table 1: $\mathbf{p = 0.001}$, $\boldsymbol{X}_0 \sim Gamma(1, 0.05)$, $\boldsymbol{X}_1 \sim Unif(0, 170)$, $\max(\boldsymbol{X}_0) = 73.0467$, $T = 138.1551$, $n_0 = n_1 = 100$, $h = (x, \log x)$, p-increment 0.0001.

Starting j	Convergence to	Iterations	
1000	0.002887173	13	Down
400	0.001487173	1	Down
300	0.001287173	1	Down
230	0.001187173	1	Down
215	0.001087173	1	Down
210	0.001087173	1	Down
200	0.001087173	1	Up
180	0.001087173	1	Up
150	0.000987172	1	Up
140	0.000987172	1	Up

A sensible estimate of $p = 0.001$ is the average from the last 7 entries which gives $\hat{p} = 0.001072887$ with absolute error of 7.2887×10^{-5}.

Table 2: $\mathbf{p = 0.0001}$, $\boldsymbol{X}_0 \sim Gamma(1, 0.05)$, $\boldsymbol{X}_1 \sim Unif(0, 210)$, $\max(\boldsymbol{X}_0) = 77.61753$, $T = 184.2068$, $n_0 = n_1 = 100$, $h = (x, \log x)$, p-increment 0.000015.

Starting j	Convergence to	Iterations	
888	0.0003439967	2	Down
577	0.0001339967	4	Down
450	0.0001189967	3	Down
350	0.0001189967	1	Down
310	0.0001039967	1	Down
300	0.0001039967	1	Down
290	0.0001039967	1	Up
280	0.0001039967	1	Up
270	0.0001039967	1	Up
260	0.0001039967	1	Up

A sensible estimate of $p = 0.0001$ is the value in the last 6 entries which gives $\hat{p} = 0.0001039967$ with absolute error of 3.9967×10^{-6}.

4.2.3 Lognormal(1,1)

Table 3: $\mathbf{p = 0.001}$, $\boldsymbol{X}_0 \sim LN(1,1)$, $\boldsymbol{X}_1 \sim Unif(0,80)$, $\max(\boldsymbol{X}_0) = 32.36495$, $T = 59.75377$, $n_0 = n_1 = 100$, $h = (x, \log x)$, p-increment 0.0001.

Starting j	Convergence to	Iterations	
1000	0.001199466	21	Down
950	0.001099466	13	Down
900	0.000999465	10	Down
800	0.000999465	5	Down
750	0.000999465	3	Down
700	0.000999465	2	Down
680	0.000999465	2	Up
680	0.000999465	2	Up
670	0.000999465	2	Up

A sensible estimate of $p = 0.001$ is the average from the last 6 entries which gives $\hat{p} = 0.000999465$ with absolute error of 5.35×10^{-7}.

Table 4: $\mathbf{p = 0.0001}$, $\boldsymbol{X}_0 \sim LN(1,1)$, $\boldsymbol{X}_1 \sim Unif(0,130)$, $\max(\boldsymbol{X}_0) = 44.82807$, $T = 112.058$, $n_0 = n_1 = 100$, $h = (x, \log x)$, p-increment 0.000015.

Starting j	Convergence to	Iterations	
800	0.0001945544	23	Down
500	0.0001795544	10	Down
300	0.0001345544	5	Down
200	0.0001195544	2	Down
170	0.0001045544	2	Down
160	0.0001045544	2	Down
155	0.0001045544	2	Up
152	0.0001045544	2	Up
150	0.0001045544	2	Up

A sensible estimate of $p = 0.0001$ is the average from the last 5 entries which gives $\hat{p} = 0.0001045544$ with absolute error of 4.5544×10^{-6}.

4.2.4 Lognormal(0,1)

Table 5: $\mathbf{p} = \mathbf{0.001}$, $\boldsymbol{X}_0 \sim LN(0,1)$, $\boldsymbol{X}_1 \sim Unif(0,50)$, $\max(\boldsymbol{X}_0) = 11.86797$, $T = 21.98218$, $n_0 = n_1 = 100$, $h = (x, \log x)$, p-increment 0.0001.

Starting j	Convergence to	Iterations	
1000	0.001099445	19	Down
900	0.001099445	5	Down
820	0.001099445	2	Down
800	0.000999444	3	Down
790	0.000999444	2	Down
780	0.000999444	2	Up
770	0.000999444	2	Up
760	0.001099445	4	Up

A sensible estimate of $p = 0.001$ is the average from the last 5 entries which gives $\hat{p} = 0.001019444$ with absolute error of 1.9444×10^{-5}.

Table 6: $\mathbf{p} = \mathbf{0.000}1$, $\boldsymbol{X}_0 \sim LN(0,1)$, $\boldsymbol{X}_1 \sim Unif(0,70)$, $\max(\boldsymbol{X}_0) = 13.77121$, $T = 41.22383$, $n_0 = n_1 = 100$, $h = (x, \log x)$, p-increment 0.000015.

Starting j	Convergence to	Iterations	
900	0.0002392241	28	Down
800	0.0001042241	25	Down
700	0.0001042241	18	Down
500	0.0001192241	6	Down
360	0.0001042241	2	Down
355	0.0001042241	2	Up
350	0.0001042241	2	Up
350	0.0001042241	2	Up

A sensible estimate of $p = 0.0001$ is the average from the last 4 entries which gives $\hat{p} = 0.0001042241$ with absolute error of 4.2241×10^{-6}.

4.2.5 f(2,7)

Table 7: $\mathbf{p = 0.001}$, $\boldsymbol{X}_0 \sim f(2,7)$, $\boldsymbol{X}_1 \sim Unif(0,50)$, $\max(\boldsymbol{X}_0) = 12.25072$, $T = 21.689$, $n_0 = n_1 = 100$, $h = (x, \log x)$, p-increment 0.0001.

Starting j	Convergence to	Iterations	
500	0.001103351	10	Down
450	0.001003351	9	Down
400	0.001003351	7	Down
300	0.001003351	4	Down
210	0.001003351	2	Up
190	0.000903350	2	Up
180	0.000903350	2	Up

A sensible estimate of $p = 0.001$ occurs at the Down-Up shift which gives $\hat{p} = 0.001003351$ with absolute error of 3.351×10^{-6}.

Table 8: $\mathbf{p = 0.0001}$, $\boldsymbol{X}_0 \sim f(2,7)$, $\boldsymbol{X}_1 \sim Unif(0,70)$, $\max(\boldsymbol{X}_0) = 14.62357$, $T = 45.13234$, $n_0 = n_1 = 100$, $h = (x, \log x)$, p-increment 0.000015.

Starting j	Convergence to	Iterations	
750	0.0001341104	3	Down
740	0.0001041104	5	Down
730	0.0001041104	4	Down
700	0.0001341104	3	Up
660	0.0001041104	2	Down
650	0.0001041104	2	Up
645	0.0001041104	2	Up
640	0.0001041104	3	Up

A sensible estimate of $p = 0.0001$ occurs at the Down-Up shift which gives $\hat{p} = 0.0001041104$ with absolute error of 4.1104×10^{-6}.

4.2.6 Weibull(0.8,2)

Table 9: $\mathbf{p = 0.001}$, $\boldsymbol{X}_0 \sim Weibull(0.8, 2)$, $\boldsymbol{X}_1 \sim Unif(0, 40)$, $\max(\boldsymbol{X}_0) = 8.081707$, $T = 22.39758$, $n_0 = n_1 = 100$, $h = (x, \log x)$, p-increment 0.0001.

Starting j	Convergence to	Iterations	
1000	0.001899263	3	Down
1000	0.001099263	8	Down
950	0.000999262	2	Immediate
950	0.000999262	2	Up
940	0.001099263	4	Up
940	0.000999262	3	Up

In the 3rd entry there was an immediate convergence. A sensible estimate of $p = 0.001$ is the average from the last 5 entries which gives $\hat{p} = 0.001039261$ with absolute error of 3.9261×10^{-5}.

Table 10: $\mathbf{p = 0.0001}$, $\boldsymbol{X}_0 \sim Weibull(0.8, 2)$, $\boldsymbol{X}_1 \sim Unif(0, 50)$, $\max(\boldsymbol{X}_0) = 12.20032$, $T = 32.09036$, $n_0 = n_1 = 100$, $h = (x, \log x)$, p-increment 0.000015.

Starting j	Convergence to	Iterations	
700	0.0002096393	21	Down
400	0.0001196393	11	Down
300	0.0001946393	2	Down
200	0.0001046393	5	Down
130	0.0001046393	2	Down
125	0.0001046393	2	Up
120	0.0001046393	2	Up
115	0.0001046393	2	Up

A sensible estimate of $p = 0.0001$ is the average from the last 5 entries which gives $\hat{p} = 0.0001046393$ with absolute error of 4.6393×10^{-6}.

4.2.7 2,4,6-trichlorophenol (ug/L)

We use trichlorophenol data from `https://wwwn.cdc.gov/nchs/nhanes` (dubbed urx3tb). There are 2604 observations of which the proportion exceeding $T = 9.5$ is $p = 0.001152074$. Consider the 2604 observations as a population and the problem is to estimate p from a sample $\boldsymbol{X}_0$ of size $n_0 = 100$ where $\max(\boldsymbol{X}_0) < T$.

The 8 estimates in Table 11 with $\max(\boldsymbol{X}_0) = 3$ seem to be in a neighborhood of the true $p = 0.001152074$. Their average is $0.001049096 \approx p$ with standard deviation of 0.5345278×10^{-5}. Note that the shift from down to up occurs at $0.001099096 \approx p$. The 3rd quartile from 10,000 B's is 0.001225, suggesting a first j around 800, as well as a reasonable first guess for p.

Repeating the iterations with different 10,000 fusions and a different $\boldsymbol{X}_0$ with $\max(\boldsymbol{X}_0) = 4.6$, we see from Table 12 that the shift from down to up

Table 11: $\mathbf{p} = \mathbf{0.001152074}$, $\boldsymbol{X}_0$ a trichlorophenol sample. $\boldsymbol{X}_1 \sim Unif(0, 30)$, $\max(\boldsymbol{X}_0) = 3$, $T = 9.5$, $n_0 = n_1 = 100$, $h = (x, \log x)$, p-increment 0.0001.

Starting j	Convergence to	Iterations	
840	0.001099096	8	Down
800	0.000999095	7	Down
760	0.000999095	4	Down
755	0.001099096	2	Down
750	0.001099096	2	Up
740	0.000999095	2	Up
735	0.000999095	2	Up
732	0.001099096	4	Up

occurs at 0.00119882 close to p with absolute error 4.6746×10^{-5}. In this case the median 0.001091 from 10,000 B's provides an approximation to p . In general, however, the 3rd quartile (here 0.003386) is a more prudent first guess.

Table 12: $\mathbf{p} = \mathbf{0.001152074}$, $\boldsymbol{X}_0$ a trichlorophenol sample. $\boldsymbol{X}_1 \sim Unif(0, 30)$, $\max(\boldsymbol{X}_0) = 4.6$, $T = 9.5$, $n_0 = n_1 = 100$, $h = (x, \log x)$, p-increment 0.0001.

Starting j	Convergence to	Iterations	
800	0.00119882	16	Down
700	0.00189882	9	Down
600	0.00119882	5	Down
590	0.00119882	6	Down
530	0.00119882	2	Up
520	0.00109882	2	Up
515	0.00109882	2	Up

4.2.8 Mercury (mg/kg)

The mercury data consists of 8,266 observations of which the proportion exceeding $T = 22.41$ is $p = 0.001088797$. The data source is NOAA's National Status and Trends Data
`https://products.coastalscience.noaa.gov/nsandt_data/data.aspx`.
The results of 10,000 fusions of a reference sample $\boldsymbol{X}_0$ with $\boldsymbol{X}_1 \sim Unif(0, 50)$ samples are summarized in Table 13. The shift from down to up occurs at 0.000999351 not far from $p = 0.001088797$ and the median from 10,000 B's is 0.001049, close to the true p.

Repeating the iterations with a different reference sample $\boldsymbol{X}_0$ as well as different 10,000 fusions with $\boldsymbol{X}_1 \sim Unif(0, 50)$ samples, we see from Table 14 that the shift from down to up occurs at 0.001099501 very close to the true $p = 0.001088797$. The median from 10,000 B's is 0.001704 giving an idea as to the magnitude of the true p.

Table 13: $\mathbf{p} = \mathbf{0.001088797}$, $\boldsymbol{X}_0$ a mercury sample. $\boldsymbol{X}_1 \sim Unif(0, 50)$, $\max(\boldsymbol{X}_0) = 7.99$, $T = 22.41$, $n_0 = n_1 = 100$, $h = (x, \log x)$, p-increment 0.0001.

Starting j	Convergence to	Iterations	
800	0.001099352	14	Down
700	0.001199352	8	Down
600	0.000999351	5	Down
500	0.000999351	2	Down
490	0.000999351	2	Up
480	0.000999351	2	Up
470	0.000999351	2	Up

Table 14: $\mathbf{p} = \mathbf{0.001088797}$, $\boldsymbol{X}_0$ a mercury sample. $\boldsymbol{X}_1 \sim Unif(0, 50)$, $\max(\boldsymbol{X}_0) = 11.9$, $T = 22.41$, $n_0 = n_1 = 100$, $h = (x, \log x)$, p-increment 0.0001.

Starting j	Convergence to	Iterations	
800	0.001199501	15	Down
700	0.001199501	12	Down
500	0.001199501	6	Down
400	0.001099501	2	Down
390	0.001099501	2	Up
380	0.001099501	2	Up
375	0.001199501	3	Up
360	0.001099501	3	Up

4.2.9 Can the Method Fail?

Problems might occur when $\max(\boldsymbol{X}_0)$ is too small or too large relative to T. A relatively small $\max(\boldsymbol{X}_0)$ indicates that the observed data are just too small, a problem that could be ameliorated by increasing $\boldsymbol{X}_0$. A large $\max(\boldsymbol{X}_0)$ indicates that the " $\bullet$ " point is too close to the lower end of the B-curve. In that case the iterative method could fail to produce converging "Up" sequences. If the reference sample is sufficiently large, the removal of few large observations creates a new smaller $\max(\boldsymbol{X}_0)$ which could push the point upward along the B-curve thus producing converging "Up" sequences.

As an example, consider the case of 10,000 fusions of $\boldsymbol{X}_0 \sim Gamma(1, 0.05)$ with $\boldsymbol{X}_1 \sim Unif(0, 160)$, where $\mathbf{p} = \mathbf{0.001}$, $T = 138.1551$, $n_0 = n_1 = 100$, $\boldsymbol{h} = (x, \log x)$, p-increment 0.0001. From Figure 3, with $\max(\boldsymbol{X}_0) = 122.1429$ the " $\bullet$ " point is at the bottom end of the B-curve slightly above $(1, B_{(1)}) = (1, 0.0008912)$. By removing the largest 3 observations from $\boldsymbol{X}_0$, the " $\bullet$ " point moves upward along the B-curve, and the iterations with the new $\max(\boldsymbol{X}_0) = 56.4284976$ and smaller sample sizes of $n_0 = n_1 = 97$ gave Down-Up sequences which converged readily to $\hat{p} = 0.00117879$. This is close to what was obtained in Table 1 with $\boldsymbol{X}_1 \sim Unif(0, 170)$; here $\boldsymbol{X}_1 \sim Unif(0, 160)$.

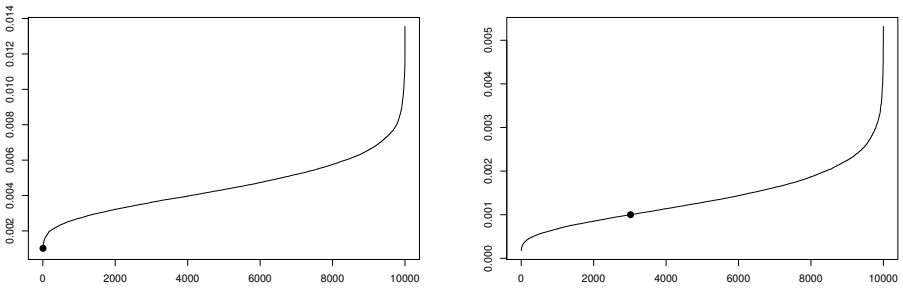

Figure 3: B-Curves, 10,000 B's, from Gamma(1,0.05), $p = 0.001$, $T = 138.1551$. Left: $n_0 = n_1 = 100$, $\max(\boldsymbol{X}_0) = 122.1429$, the point "•" is at the bottom of the B-curve slightly above $(1, B_{(1)}) = (1, 0.0008912)$. Right: Three largest observations removed, $n_0 = n_1 = 97$, $\max(\boldsymbol{X}_0) = 56.4284976$, the point "•" moves upward along the B-curve.

5 Variability of Point Estimates

Clearly, the iterative method (IM) can be repeated many times with different B-samples of size 1,000 taken from, say, 10,000 B's (10,000 is our default) to produce tail probability estimates as above from which variance approximations can be obtained. The following tables show typical variance approximations, obtained from single convergent sequences where the starting j corresponds to the 3rd quartile of the sampled 1,000 B's. Each entry was obtained from 1,000 runs for both $n_0 = n_1 = 100$ and $n_0 = n_1 = 200$. There is improvement in precision going from samples of size 100 to 200. In all cases reported here, and many other additional cases, $\sigma_{\hat{p}} = \mathrm{O}(10^{-4})$. In the tables $\bar{p}$ is the average estimate of p from 1,000 runs.

Table 15: $X_0 \sim$ **2,4,6-trichlorophenol (Urx3tb)**: p = 1 - G(T) = 0.001152074, T = 9.50, $X_1 \sim$ Unif(0,12), $n_0 = n_1, h(x) = (x, \log x)$.

	$\bar{p}$	$\sigma_{\hat{p}}$
$n_0 = 100, \max(\boldsymbol{X}_0) = 8.8$	0.0011539	0.0004269
$n_0 = 200, \max(\boldsymbol{X}_0) = 8.8$	0.0011216	0.0002871

Table 16: $X_0 \sim \mathbf{LN}(0,1)$: $p = 1 - G(T) = 0.001, T = 21.98218, X_1 \sim$ Unif(1,60), $n_0 = n_1, h(x) = (x, \log x)$.

	$\bar{p}$	$\sigma_{\hat{p}}$
$n_0 = 100, \max(\boldsymbol{X}_0) = 11.04102$	0.0011401	0.0004100
$n_0 = 200, \max(\boldsymbol{X}_0) = 11.04102$	0.0010598	0.0002823

Table 17: $X_0 \sim$ **Weibull**$(1,2)$: $p = 1 - G(T) = 0.001, T = 13.81551, X_1 \sim$ Unif(0,16), $n_0 = n_1, h(x) = (x, \log x)$.

	$\bar{p}$	$\sigma_{\hat{p}}$
$n_0 = 100, \max(\boldsymbol{X}_0) = 8.626444$	0.0011215	0.0001524
$n_0 = 200, \max(\boldsymbol{X}_0) = 8.673713$	0.0010768	0.0001372

Table 18: $X_0 \sim$ **Pareto**$(1,4)$: $p = 1 - G(T) = 0.001, T = 5.623413, X_1 \sim$ Unif(1,8), $n_0 = n_1, h(x) = (x, \log x)$.

	$\bar{p}$	$\sigma_{\hat{p}}$
$n_0 = 100, \max(\boldsymbol{X}_0) = 3.08099$	0.0011549	0.0002256
$n_0 = 200, \max(\boldsymbol{X}_0) = 4.14516$	0.0009966	0.0002034

6 Comparison: ROFS vs POT

From a practical view point, some comparison is needed to assess the relative merit of ROSF/IM. We provide in what follows a limited comparison against a well known method, however, a more extensive comparison is warranted and will be dealt with elsewhere.

Thus, against the background provided in the previous sections, we compare two very different ways to obtain interval estimates for small tail probabilities. The well known peaks over threshold (POT) based on extreme value theory, and the present iterative process based on repeated fusion of a given reference sample with external computer-generated uniformly distributed samples. The comparison is based on confidence interval coverage, interval width, and on the mean absolute error (MAE) which measures the discrepancy between $\hat{p}$ and the true tail probability p. In Tables 19 to 27, p is relatively small, $p = 0.001$ (or approximately so), whereas in the last two Tables 28 and 29, p is smaller, $p = 0.0001$.

Throughout the comparison the sample sizes are $n_0 = n_1 = 100$ or $n_0 = n_1 = 200$, and $\boldsymbol{h}(x) = (x, \log x)$. Thus, in the present comparison the reference $\boldsymbol{X}_0$ and the fusion samples $\boldsymbol{X}_1$ have size $n_0 = 100$ or $n_0 = 200$.

To save computation time, in each case of the iteration process F_B was obtained from 1000 fusions, and we use in each case a single convergent sequence where the starting j is such that $B_{(j)}$ is approximately equal to the 3rd quartile of the observed 1000 B's. Starting at the 3rd quartile is computationally sensible as the corresponding $B_{(j)}$ most often converge to a point in a neighborhood of p as j increases. See Tables 3 to 14 and more examples in Kedem et al. (2018).
The following tables are the result of 500 runs. In each run the iteration method (IM) was repeated 500 times.

From the mean residual life (MRL) plots we obtained the thresholds u needed for the POT method. In all cases reported in the tables, the MRL plots suggest the use of the largest 20% of the reference data $\boldsymbol{X}_0$ for fitting the generalized Pareto (GP) distribution. We have noticed a deterioration in the POT results when using 30%, 15% or 10% of $\boldsymbol{X}_0$. The simulation details are given in Section A.2 in the Appendix.

An interesting picture emerges from Tables 19 to 29. For moderately large sample sizes of $n_0 = 100$ and $n_0 = 200$, regardless of the tail type, already with $\mathbf{N = 50}$, that is, the number of $\hat{p}$'s used in forming the CI for the true p of the form $(\min(\hat{p}), \max(\hat{p}))$ (defined in Section A.2), the iteration process gives reliable and relatively narrow confidence intervals, whereas the POT gives unacceptable coverage and in most cases wider CI's and greater MAE as well. The POT coverage increases significantly going from $n_0 = 100$ to $n_0 = 200$, however, it seems that for the method to "fire up" larger samples are needed. Regarding ROSF, the choice of $\mathbf{N = 50}$ seems prudent across all cases, and with $n_0 = 200$ shorter CI's achieve coverage similar to that from the smaller $n_0 = 100$. In all cases the MAE from the iteration process is much smaller than that obtained from POT.

6.1 Comparison Tables

The following tables compare ROSF and POT for $p = 0.001$ and $p = 0.0001$.

Table 19: $X_0 \sim \mathbf{t}_{(1)} > 0 : p = 1 - G(T) = 0.001, T = 631.8645, X_1 \sim$ Unif(0,800), $n_0 = n_1, \boldsymbol{h}(x) = (x, \log x)$. p-increment 0.0001.

		$n_0 = 100$			$n_0 = 200$		
Method	N	Coverage	CI Length	MAE	Coverage	CI Length	MAE
POT	-	63.2%	0.00372	0.00149	72.1%	0.00292	0.00122
ROSF & IM	**50**	98.2%	0.00213	0.00061	100%	0.00193	0.00051
	100	100%	0.00264	-	100%	0.00241	-

Table 20: $X_0 \sim \mathbf{Weibull}(1, 2) : p = 1 - G(T) = 0.001, T = 13.81551, X_1 \sim$ Unif(0,16), $n_0 = n_1, \boldsymbol{h}(x) = (x, \log x)$. p-increment 0.00005.

		$n_0 = 100$			$n_0 = 200$		
Method	N	Coverage	CI Length	MAE	Coverage	CI Length	MAE
POT	-	82.7%	0.00431	0.00131	87.8%	0.00333	0.00083
ROSF & IM	**50**	92.5%	0.00287	0.00068	92.8%	0.00231	-
	100	100%	0.00381	-	100%	0.00321	-

Table 21: $X_0 \sim \mathbf{Pareto}(1,4) : p = 1 - G(T) = 0.001, T = 5.623413, X_1 \sim$ Unif(1,8), $n_0 = n_1, \boldsymbol{h}(x) = (x, \log x)$. p-increment 0.0001.

		$n_0 = 100$			$n_0 = 200$		
Method	N	Coverage	CI Length	MAE	Coverage	CI Length	MAE
POT	-	81.8%	0.00419	0.00121	84.5%	0.00337	0.00070
ROSF & IM	**50**	96.2%	0.00232	0.00052	97.8%	0.00231	0.00041
	100	100%	0.00272	-	100%	0.00269	-

Table 22: $X_0 \sim \mathbf{Gamma}(3,1) : p = 1 - G(T) = 0.001, T = 11.22887, X_1 \sim$ Unif(0,20), $n_0 = n_1, \boldsymbol{h}(x) = (x, \log x)$. p-increment 0.00005.

		$n_0 = 100$			$n_0 = 200$		
Method	N	Coverage	CI Length	MAE	Coverage	CI Length	MAE
POT	-	77.3%	0.00410	0.00096	86.1%	0.00321	0.00081
ROSF & IM	**50**	93.4%	0.00188	0.00054	94.5%	0.00175	0.00043
	100	100%	0.00256	-	100%	0.00248	-

Table 23: $X_0 \sim \mathbf{IG}(2,40) : p = 1 - G(T) = 0.001, T = 3.835791, X_1 \sim$ Unif(0,8), $n_0 = n_1, \boldsymbol{h}(x) = (x, \log x)$. p-increment 0.00005.

		$n_0 = 100$			$n_0 = 200$		
Method	N	Coverage	CI Length	MAE	Coverage	CI Length	MAE
POT	-	69.6%	0.00324	0.00123	82.3%	0.00316	0.00092
ROSF & IM	**50**	100%	0.00289	0.00047	100%	0.00206	0.00041
	100	100%	0.00332	-	100%	0.00313	-

Table 24: $X_0 \sim \mathbf{LN}(0,1) : p = 1 - G(T) = 0.001, T = 21.98218, X_1 \sim$ Unif(1,60), $n_0 = n_1, \boldsymbol{h}(x) = (x, \log x)$. p-increment 0.00005.

		$n_0 = 100$			$n_0 = 200$		
Method	N	Coverage	CI Length	MAE	Coverage	CI Length	MAE
POT	-	81.5%	0.00451	0.00111	85.2%	0.00392	0.00103
ROSF & IM	**50**	100%	0.00234	0.00047	100%	0.00199	0.00039
	100	100%	0.00267	-	100%	0.00244	-

Table 25: $X_0 \sim \mathbf{LN}(1,1) : p = 1 - G(T) = 0.001, T = 59.75377, X_1 \sim$ Unif(1,140), $n_0 = n_1, \boldsymbol{h}(x) = (x, \log x)$. p-increment 0.0001.

		$n_0 = 100$			$n_0 = 200$		
Method	N	Coverage	CI Length	MAE	Coverage	CI Length	MAE
POT	-	81.4%	0.00435	0.00117	86.8%	0.00399	0.00099
ROSF & IM	**50**	89.1%	0.00187	0.00069	100%	0.00164	0.00052
	100	100%	0.00199	-	100%	0.00192	-

Table 26: $X_0 \sim$ Mercury : $p = 1 - G(T) = 0.001088797, T = 22.41, X_1 \sim$ Unif(0,50), $n_0 = n_1, \boldsymbol{h}(x) = (x, \log x)$. p-increment 0.0001.

		$n_0 = 100$			$n_0 = 200$		
Method	N	Coverage	CI Length	MAE	Coverage	CI Length	MAE
POT	-	85.3%	0.00455	0.00130	88.6%	0.00398	0.00122
ROSF & IM	**50**	97.5%	0.00215	0.00048	100%	0.00197	0.00045
	100	100%	0.00259	-	100%	0.00238	-

Table 27: $X_0 \sim$ URX3TB : $p = 1 - G(T) = 0.001152074, T = 9.50, X_1 \sim$ Unif(0,12), $n_0 = n_1, \boldsymbol{h}(x) = (x, \log x)$. p-increment 0.0001. Data source for URX3TB - 2,4,6-trichlorophenol (ug/L): `https://wwwn.cdc.gov/nchs/nhanes`

		$n_0 = 100$			$n_0 = 200$		
Method	N	Coverage	CI Length	MAE	Coverage	CI Length	MAE
POT	-	81.1%	0.00433	0.00143	87.1%	0.00376	0.00123
ROSF & IM	**50**	89.1%	0.00179	0.00055	96.9%	0.00177	0.00044
	100	100%	0.00241	-	100%	0.00235	-

Table 28: $X_0 \sim \mathbf{F}(2, 12) : p = 1 - G(T) = 0.0001, T = 21.84953, X_1 \sim$ Unif(0,25), $n_0 = n_1, \boldsymbol{h}(x) = (x, \log x)$. p-increment 0.00001.

		$n_0 = 100$			$n_0 = 200$		
Method	N	Coverage	CI Length	MAE	Coverage	CI Length	MAE
POT	-	71.4%	0.00062	0.00052	81.6%	0.00053	0.000045
ROSF & IM	**50**	95.2%	0.00059	0.00022	96.3%	0.00052	0.000019
	100	100%	0.00082	-	100%	0.00069	-

Table 29: $X_0 \sim \mathbf{LN}(0, 1) : p = 1 - G(T) = 0.0001, T = 41.22383, X_1 \sim$ Unif(1,60), $n_0 = n_1, \boldsymbol{h}(x) = (x, \log x)$. p-increment 0.00001.

		$n_0 = 100$			$n_0 = 200$		
Method	N	Coverage	CI Length	MAE	Coverage	CI Length	MAE
POT	-	72.1%	0.00064	0.00045	82.6%	0.00047	0.000039
ROSF & IM	**50**	100%	0.00066	0.00021	100%	0.00057	0.000017
	100	100%	0.00083	-	100%	0.00079	-

7 Discussion

The numerous number of fusions of a given reference sample with computer generated samples gives rise to different observables including the upper bounds for a tail probability p that were used in the paper. The upper bounds, obtained from the combined real and artificial data, were mostly much larger than p, some were less than p, but some among the multitude of upper bounds essentially coincided with p and they were identified rather closely using an iterative procedure.

We have illustrated that, across a fairly wide range of distributional tail types, repeated fusion of a reference sample with externally generated uniform random data allowed us to gain information about the tail behavior beyond the threshold using the notion of B-curves coupled with a well known formula from order statistics. In neighborhoods of the true p, the consequential Down-Up sequences tended to transition or shift at points close to p, providing surprisingly close estimates. We have seen that with sample sizes on the order of 100 we can in many cases estimate tail probabilities on the order of 1/10,000. It seems that larger samples are needed for much smaller tail probabilities, and that the method could fail when $\max(\boldsymbol{X}_0)$ is exceedingly small or exceedingly large relative to the threshold T. That is, when the "$\bullet$" point on the B-curve is very close to one of the two ends of the curve.

Throughout the paper the fusion samples were uniform samples whose support contained T. That is, the upper limit of the support exceeded T. But other than this, no guide for choosing the upper limits was provided. Experience, however, shows that different upper limits give similar results.

The ideas presented in this paper can be extended in a number of ways. For example, using "fake" data from distributions other than uniform, and using different fusion mechanisms together with appropriate inferential methods other than the semiparametric method used in the paper. That is, explore different ways of connecting $\boldsymbol{X}_0$ and $\boldsymbol{X}_1$, other than by means of their distributions as expressed by the density ratio model.

Reliable estimation of tail probabilities is important in numerous fields from finance to geophysics to meteorology to the design of ships and to optics; see Pelinovsky et al. (2008), and Solli et al. (2007).

A Appendix

The appendix addresses the density ratio model (5) for $m+1$ data sources discussed briefly in Section 3.

A.1 Asymptotic Distribution of $\hat{G}(x)$

Define $\alpha_0 \equiv 0, \beta_0 \equiv 0$, $w_j(x) = \exp(\alpha_j + \beta_j' h(x))$, $\rho_i = n_i/n_0$, $j = 1, \ldots, m$.

Maximum likelihood estimates for all the parameters and $G(x)$ can be obtained by maximizing the empirical likelihood over the class of step cumulative distribution functions with jumps at the observed values $t_1, \ldots, t_n$ (Owen 2001). Let $p_i = dG(t_i)$ be the mass at t_i, for $i = 1, \ldots, n$. Then the empirical likelihood becomes

$$\mathcal{L}(\boldsymbol{\theta}, G) = \prod_{i=1}^{n} p_i \prod_{j=1}^{n_1} \exp(\alpha_1 + \beta_1' h(x_{1j})) \cdots \prod_{j=1}^{n_m} \exp(\alpha_m + \beta_m' h(x_{mj})). \quad (10)$$

Maximizing $\mathcal{L}(\boldsymbol{\theta}, G)$ subject to the constraints

$$\sum_{i=1}^{n} p_i = 1, \ \sum_{i=1}^{n} p_i[w_1(t_i) - 1] = 0, \ldots, \sum_{i=1}^{n} p_i[w_m(t_i) - 1] = 0 \quad (11)$$

we obtain the desired estimates. In particular,

$$\hat{G}(t) = \frac{1}{n_0} \cdot \sum_{i=1}^{n} \frac{I(t_i \le t)}{1 + \rho_1 \exp(\hat{\alpha}_1 + \hat{\beta}_1' h(t_i)) + \cdots + \rho_m \exp(\hat{\alpha}_m + \hat{\beta}_m' h(t_i))}, \quad (12)$$

where $I(t_i \le t)$ equals one for $t_i \le t$ and is zero, otherwise. Similarly, $\hat{G}_j$ is estimated by summing $\exp(\hat{\alpha}_j + \hat{\beta}_j' h(t_i)) dG(t_i)$.

The asymptotic properties of the estimators have been studied by a number of authors including Qin and Zhang (1997), Lu (2007), and Zhang (2000).

Define the following quantities: $\boldsymbol{\rho} = \text{diag}\{\rho_1, \ldots, \rho_m\}$,

$$A_j(t) = \int \frac{w_j(y) I(y \le t)}{\sum_{k=0}^{m} \rho_k w_k(y)} dG(y), \quad B_j(t) = \int \frac{w_j(y) h(y) I(y \le t)}{\sum_{k=0}^{m} \rho_k w_k(y)} dG(y),$$

$$\bar{A}(t) = (A_1(t), \ldots, A_m(t))', \quad \bar{B}(t) = (B_1'(t), \ldots, B_m'(t))'.$$

Then the asymptotic distribution of $\hat{G}(t)$ for $m \ge 1$ is given by the following result due to Lu (2007).

Theorem A.1 *Assume that the sample size ratios $\rho_j = n_j/n_0$ are positive and finite and remain fixed as the total sample size $n = \sum_{j=0}^{m} n_j \to \infty$. The process $\sqrt{n}(\hat{G}(t) - G(t))$ converges to a zero-mean Gaussian process in the space of real right continuous functions that have left limits with covariance matrix given by*

$$\begin{aligned}\mathrm{Cov}\{\sqrt{n}(\hat{G}(t) - G(t)), \sqrt{n}(\hat{G}(s) - G(s))\} = \\ \left(\sum_{k=0}^{m} \rho_k\right)\left(G(t \wedge s) - G(t)G(s) - \sum_{j=1}^{m} \rho_j A_j(t \wedge s)\right) \\ + \left(\bar{A}'(s)\boldsymbol{\rho}, \bar{B}'(s)(\boldsymbol{\rho} \otimes I_p)\right) S^{-1} \begin{pmatrix} \boldsymbol{\rho}\bar{A}(t) \\ (\boldsymbol{\rho} \otimes I_p)\bar{B}(t) \end{pmatrix}. \qquad (13)\end{aligned}$$

where I_p is the $p \times p$ identity matrix, and $\otimes$ denotes Kronecker product.

For a complete proof see Lu (2007). The proof for $m = 1$ is given in Zhang (2000).

Denote by $\hat{V}(t)$ the estimated variance of $\hat{G}(t)$ as given in (13). Replacing parameters by their estimates, a $1-\alpha$ level pointwise confidence interval for $G(t)$ is approximated by

$$\left(\hat{G}(t) - z_{\alpha/2}\sqrt{\hat{V}(t)}, \;\; \hat{G}(t) + z_{\alpha/2}\sqrt{\hat{V}(t)}\right), \qquad (14)$$

where $z_{\alpha/2}$ is the upper $\alpha/2$ point of the standard normal distribution. Hence, a $1-\alpha$ level pointwise confidence interval for $1 - G(T)$ for any T, and in particular for relatively large thresholds T is approximated by

$$\left(1 - \hat{G}(t) - z_{\alpha/2}\sqrt{\hat{V}(t)}, 1 - \hat{G}(t) + z_{\alpha/2}\sqrt{\hat{V}(t)}\right). \qquad (15)$$

A.2 Simulation Description

The following steps were followed. There were 500 runs. In each run the iteration method (IM) was repeated 500 times.

First, a reference $\boldsymbol{X}_0$ was obtained.

POT:
The POT procedure was applied to get both an estimate $\hat{p}$ and a confidence

interval (CI). The MRL plots suggest the use of the largest 20% of the reference data $\boldsymbol{X}_0$ for fitting the generalized Pareto (GP) distribution.

ROSF/IM:
$\boldsymbol{X}_0$ was fused with $\boldsymbol{X}_1$ 1000 times (ROSF) to get F_B and then $\hat{p}$ (IM).
$\boldsymbol{X}_0$ was fused again with different $\boldsymbol{X}_1$ 1000 times to get F_B and $\hat{p}$.
This was repeated 500 times.

The iterative method thus gave 500 $\hat{p}$'s. We then chose at random N $\hat{p}$'s from 500 $\hat{p}$'s to construct a CI for the true p as $(\min(\hat{p}), \max(\hat{p}))$.
This is run 1.

The above steps were repeated, for both POT and ROSF/IM each time with a different $\boldsymbol{X}_0$, 500 times (runs) to obtain coverage and average CI length. In the tables, CI length is an average length from 500 intervals.

Since there are 500 runs, POT gave 500 $\hat{p}$'s. Regarding IM, a single $\hat{p}$ was chosen at random (out of 500 $\hat{p}$'s) from each of the 500 runs. The mean absolute error (MAE) was obtained in both cases from the mean of 500 absolute differences $\sum(|\hat{p}_i - p|)/500$, where $p = 0.001$ or $p = 0.0001$. In the iterative method, in each table the MAE is reported once on the line corresponding to $N = 50$.

References

[1] Beirlant, J., Goegebeur, Y., Teugels, J.L., and Segers, J. *Statistics of extremes : theory and applications.* Wiley: Hoboken, 2004.

[2] Ferreira, A. and De Haan, L. On the block maxima method in extreme value theory: PWM estimators. *The Annals of Statistics* 2015, **43**: 276-298.

[3] Fithian, W. and Wager, S. Semiparametric exponential families for heavy-tailed data. *Biometrika* 2015, **102**: 486-493.

[4] Fokianos, K. and Qin J. A Note on Monte Carlo Maximization by the Density Ratio Model. *Journal of Statistical Theory and Practice* 2008; **2**: 355-367.

[5] Katzoff, M., Zhou, W., Khan, D., Lu, G., and Kedem, B. Out of sample fusion in risk prediction. *Journal of Statitsical Theory and Practice* 2014; **8**: 444-459.

[6] Kedem, B., Pan, L., Zhou, W., and Coelho, C.A. Interval estimation of small tail probabilities – application in food safety. *Statistics in Medicine* 2016; **35**: 3229-3240.

[7] Kedem, B., Pan, L., Smith, P. and Wang, C. (2018). Repeated out of sample fusion in the estimation of small tail probabilities. arXiv:1803.10766v2 [stat.ME], May 2018.

[8] Kedem, B., De Oliveira, V., and Sverchkov, M. (2017). *Statistical Data Fusion*, World Scientific, Singapore.

[9] Lu, G. Asymptotic Theory for Multiple-Sample Semiparametric Density Ratio Models and its Application to Mortality Forecasting. Ph.D. Dissertation, University of Maryland, College Park, 2007.

[10] Pelinovsky, E. and Kharif, C., Editors. *Extreme Ocean Waves*, Springer, New York, 2008

[11] Qin, J. and Zhang, B. A Goodness of fit test for logistic regression models based on case-control data. *Biometrika* 1997; **84**: 609-618.

[12] Solli, D.R., Ropers, C., Koonath, P., and Jalali, B. Optical rogue waves. *Nature* 2007; **450**: 1054-1057.

[13] Zhang, B. A goodness of fit test for multiplicative-intercept risk models based on case-control data. *Statistica Sinica* 2000; **10**: 839-865.

[14] Zhou, W. Out of Sample Fusion. Ph.D. Dissertation, University of Maryland, College Park, 2013.