TECHNOMETRICS©, VOL. 16, NO. 1, FEBRUARY 1974

Heavy-Tailed Distributions: Properties and Tests

Maurice C. Bryson

Department of Statistics
Colorado State University
Fort Collins, Colorado

Distributions with heavier-than-exponential tails are studied for describing empirical phenomena. It is argued that the concept of increasing "conditional mean exceedance" provides a reasonable way of describing the heavy-tail phenomenon, and a family of Pareto distributions is shown to represent distributions for which this parameter is linearly increasing. A test is developed and modified so as to be suitable for testing heavy-tailedness, and some graphical procedures are also suggested.

KEY WORDS

Heavy-tailed Distributions
Distributions
Hypothesis Tests
Pareto Distributions
Mean Residual Lifetime

1. INTRODUCTION

Despite the wide variety of continuous statistical distributions encountered in theoretical work only a relatively few families of them are commonly used in applications as models for empirical phenomena. Classically, of course, the normal distribution has nearly two centuries of use as a "law of errors" for the empirical data. The widespread use of the exponential family, particularly for lifetime data, seems to date only since the work of Davis (1952) and Epstein and Sobel (1953). Barlow and Proschan (1965) note the popularity of the gamma, Weibull, lognormal, and modified extreme value families as alternatives for lifetime distribution models. Since the work of Gumbel (1958), the extreme-value distribution has been important in hydrological and meteorological study, but for describing empirical data the lognormal and the gamma family (including exponential) seem to be by far the most widely used; see, for example, Yevjevich (1971). The gamma family also shows up in operations research applications under its alias of the Erlang distribution (see Wagner (1969)). Finally, the Pareto family has found application only as a distribution for income statistics and a few other specialized variables (see Johnson and Kotz (1970) and Malik (1970)). These and other families have been conveniently summarized and extensively discussed by Johnson and Kotz.

The experimenter today, with reference to a long history of use for one particular family of distributions in his area, typically makes an a priori selection of the family and then estimates its parameters in order to get a best fit. It is in the selection of the family, though, that an important mistake can be made. In particular, it is difficult to make a selection of a family that will suitably model the "tails" of the distribution. By definition, there will not be many observations from the tails of the distribution to be fit, so that a chi-square goodness-of-fit test can never reveal an ill-fitting tail without a very large amount of data. Likewise, the empirical *cdf* will be close to zero or one in the tails, so that a Kolmogorov–Smirnov test will also fail.

Nonetheless, failure to suitably fit the tail of the distribution can be very serious even if the fitted distribution "looks good" in terms of the resemblance between its density function and a frequency histogram. Suppose, for example, that a set of flood data were erroneously modeled using a one-sided normal distribution when in fact a one-sided Cauchy distribution was called for. As every student notices in his first course in distribution theory, the normal and Cauchy density functions look remarkably similar. However, the hypothesized error would result in underestimating the magnitude of a 100-year flood, compared to that of a 10-year flood, by a factor of 6.44; for the 200-year flood, he would similarly underestimate by a factor of 11.8.

This example is of course artificial, particularly so in that the experimenter wouldn't be using a normal distribution; ordinarily he would be using a gamma. But the same principle applies, and it is therefore that we should consider the question: when is a gamma tail too light? Henceforth we will in this paper restrict the problem to one of fitting continuous non-negative data (such as hydrologic, meteorological, or life testing data), and concentrate on the question of whether the actual tail of the distribution is heavier than a gamma tail. The gamma tail is, of course, an exponential tail since the polynomial component becomes relatively un-

Received Nov. 1972; revised April 1973

important at the extreme. Based on this empirical situation, then, we may say without being excessively arbitrary that "heavy-tailed" means having a density function that goes to zero less rapidly than an exponential function. As will be seen, this is a convenient definition for theoretical reasons as well. It will be noted that most of the commonly used distributions mentioned above have either exponential or lighter-than-exponential tails; the heavy-tailed exceptions are the lognormal and Pareto families, and rarely used sub-families of the Weibull.

There is evidence that heavy-tailed distributions may in fact occur more commonly than would be supposed from the infrequency of their use. Mielke (1973), in investigating precipitation data, found that they could better be fit by "kappa" distributions, a heavy-tailed family with distribution function

$$F(x) = [x^{\alpha\theta}/(\alpha + x^{\alpha\theta})]^{1/\alpha}$$

than with gamma distributions. Granger and Orr (1972) suggest the use of infinite-variance stable distributions as possible fits for various types of economic data. The stable distributions also seem to be important in hydrology; Boes and Salas–La Cruz (1973) have shown them to be of great value in explaining the "Hurst phenomenon" of partial sum ranges that increase faster than the square root of sample size.

2. The Conditional Mean Exceedance (CME)

If one were to approach the subject of heavy-tailedness from a theoretical viewpoint, he might argue as follows: a distribution has a "heavy tail" if there tend to be many large exceedances of a given magnitude. "Many" and "large" are comparative terms, so he might be more specific by saying that the average exceedance (the concept of average incorporating both the number and the magnitude of the exceedance) tends to get bigger as you get farther out into the tail of the distribution. This leads immediately to the concept of the *conditional mean exceedance*, which is here defined as

$$\mathrm{CME}_x = E(X - x \mid X \geq x).$$

In words, CME_x is the average amount by which the random variable X exceeds a given x, conditional on that exceedance being non-negative. If X is a lifetime, as in demographic or reliability work, CME_x will be recognized as the "mean residual lifetime".

This line of argument, then, leads to the definition of a heavytailed distribution as one for which CME_x is an increasing function of x—at least for sufficiently large x, since our concern is only with the tail of the distribution. Likewise, a decreasing CME_x characterizes a light-tailed distribution, and a distribution for which CME_x is constant is the borderline case.

But this is immediately seen to be the same definition as that arrived at empirically, for from the definition

$$\mathrm{CME}_x = \frac{1}{\bar{F}(x)} \int_x^\infty \bar{F}(t)\, dt$$

where $\bar{F}(x) = 1 - F(x)$ is the tail probability function. Solving the equation for constant CME_x shows that if $\mathrm{CME}_x = 1/\lambda$, then the distribution is exponential with rate parameter λ.

Thus, we may characterize the class of heavy-tailed distributions as those for which the CME is, for large x, an increasing function of x. This will henceforth be denoted as ICME (increasing conditional mean exceedance). It may be seen from the above formula that CME_x depends only on the values of $\bar{F}(t)$ for $t \geq x$, so that the CME does in fact provide a representation of the tail behavior as was desired.

3. Testing for ICME

Because of its interpretation as mean residual lifetime, the CME has been studied primarily in connection with demographic and reliability statistics. Bryson and Siddiqui (1969) identified *decreasing* mean residual lifetime as one of seven criteria that may be used to describe the "aging" of an entity, and a test for decreasing mean residual lifetime was suggested by Bryson (1968). It was previously shown by Barlow and Proschan (1965) that decreasing mean residual lifetime is implied by the more commonly used "aging criterion" of increasing failure rate (or hazard rate). But the mean residual lifetime is a weaker parameter than the hazard rate, in that its behavior depends only on the tail of the distribution. Many authors have considered the question of testing for increases or decreases in the hazard rate, as summarized by Fercho and Ringer (1972). Hollander and Proschan (1972) developed a related test for a "new better than used" (NBU) criterion, but this criterion depends primarily on the behavior of the lifetime distribution for small x rather than for large x. Among the various descriptors of lifetime distributions that can be used, the mean residual lifetime seems to have had relatively little attention.

We now consider testing the null hypothesis CME_x = constant against the alternative that it is increasing. To be more specific, a reasonable choice of alternatives is the set of linear functions

$$\mathrm{CME}_x = a + bx.$$

Applying the definition of CME_x and solving the resultant differential equation yields immediately

$$\bar{F}(x) = \left(\frac{a}{a + bx}\right)^{1+1/b}$$

as the alternative hypothesis ($b > 0$). Letting $b \to 0$ gives as the limiting case

$$\bar{F}(x) = e^{-x/a}$$

as intended. It is noted by Johnson and Kotz that the alternative distributions are sometimes called the family of Lomax distributions, or Pareto distributions of the second kind. Although originally suggested by Pareto as one model for income distributions, the family does not seem to have been much in use this century.

In considering possible tests, we note first that there is a natural condition of invariance that applies, in that the test should be independent of the measurement scale for X. Thus we impose the constraint that the test be invariant with respect to the group of transformations

$$g_k(x) = kx, \qquad k > 0.$$

With respect to this group, a maximal invariant for the sample $(x_1, x_2, \cdots, x_n)$ is $(y_1, y_2, \cdots, y_{n-1})$ where

$$y_i = x_{(i)}/x_{(n)} \qquad i = 1, 2, \cdots, n-1$$

and the ordered sample is

$$x_{(1)} < x_{(2)} < \cdots < x_{(n)} .$$

Equivalent y samples must obviously be derived from data which differ by no more than a scale factor, so any invariant test must be based on the y statistics. Application of routine transformation mechanics gives as the distribution of the y statistics under the null hypothesis,

$$f_v(y_1, y_2, \cdots, y_{n-1}) = \frac{n!\,(n-1)!}{\omega^n}$$

where

$$\omega = 1 + \sum_1^{n-1} y_i = \sum_{i=1}^{n} x_i/x_{(n)} ,$$

and the density is over the range

$$0 < y_1 < y_2 < \cdots < y_{n-1} < 1.$$

Under the alternative hypothesis, a similar procedure yields the result

$$f_v(y_1, y_2, \cdots, y_{n-1}) = n!\,(1 + b)^n \cdot \int_0^\infty u^{n-1}\left[\frac{1}{1+u}\prod_1^{n-1}\frac{1}{1+y_i u}\right]^{(2b+1)/b} du,$$

over the same range. The expression, unfortunately, remains intractable even in the simplest of special cases.

Without sacrificing the basic approach of the problem, we can simplify the mathematics considerably by the following device. If X does have the indicated Lomax distribution, then

$$Z = X + A$$

will have the more tractable Pareto distribution

$$\left.\begin{aligned}\bar{F}(z) &= 1 \qquad & 0 < z \le A\\ &= \left(\frac{A}{z}\right)^K & z > A\end{aligned}\right\} A, K > 0$$

where, in terms of the previous notation,

$$K = 1 + 1/b$$

and

$$A = a/b.$$

(Note that while the distribution exists for all $K > 0$, the mean exists only when $K > 1$.)

A scale-invariant test can now be developed on the basis of the statistics $z_1, \cdots, z_n$. The price to be paid for the simplification of the mathematics is that it will be necessary to estimate the parameter A before the x-statistics can be used. The density function of Z is now

$$f(z) = (K/A)(A/z)^{K+1} \qquad (z > A).$$

For a scale-invariant test, take $A = 1$ without loss of generality. Then, transform to the maximal invariant as before, with $y_i = z_{(i)}/z_{(n)}$:

$$f(y_1, y_2, \cdots, y_{n-1}) = (n-1)!\, K^{n-1}(y_1)^{nK}\left(\prod_{i=1}^{n-1} y_i\right)^{-K-1},$$

again over the range

$$0 < y_1 < y_2 < \cdots < y_{n-1} < 1.$$

Substituting for the y's gives a likelihood function proportional to

$$[(x_{(1)} + A)/(x_{(n)} + A)]^{nK} \cdot \left[\prod_{i=1}^{n} \{(x_{(i)} + A)/(x_{(n)} + A)\}\right]^{-K-1};$$

this may be written

$$\left[\frac{x_{(1)} + A}{\bar{x}_{GA}}\right]^{nK+n} \cdot \left[\frac{x_{(n)} + A}{x_{(1)} + A}\right]^n,$$

where $\bar{x}_{GA}$ denotes the geometric mean modified by A,

$$\bar{x}_{GA} = \left(\prod_{i=1}^{n} (x_i + A)\right)^{1/n}$$

Dividing by the likelihood under the null hypothesis gives a likelihood ratio proportional to

$$\left[\frac{x_{(n)} + A}{x_{(1)} + A}\right]^{n}\left[\frac{\bar{x}}{x_{(n)}}\right]^{n}\left[\frac{x_{(1)} + A}{\bar{x}_{GA}}\right]^{nK+n}$$

Again it is evident that no uniformly best test can be found. However, the following procedure appears to be a reasonable one. First, select a convenient value of K for which the most-powerful test could be selected. Such a value would be $K = 1$, which represents the limiting case of distributions for which the mean exists, and as such is very heavy-tailed (tail behavior is the same as for the Cauchy distribution). Second, choose A so that the theoretical *cdf* matches the empirical *cdf* at the point represented by the largest observation. This estimation procedure guarantees that the tail behavior will be modeled well, though possibly at the expense of other parts of the distribution. This constraint gives

$$\hat{A} = x_{(n)}/(n^{1/K} - 1),$$

or for the case $K = 1$, simply $x_{(n)}/(n - 1)$. After some simplification, the likelihood ratio is now proportional to

$$T = \frac{\bar{x}\left(x_{(1)} + \dfrac{x_{(n)}}{n - 1}\right)}{\bar{x}_{GA}^{\ 2}}.$$

T could be used as a test statistic except for one remaining problem. Under the assumption that X has an exponential distribution, $X_{(1)}$ will usually be small compared to $x_{(n)}/n - 1$, but this will not necessarily be true if X has, say, a gamma distribution with a large shape parameter. Accordingly the use of T might run the risk of rejecting the null hypothesis because of a large $x_{(1)}$ rather than because of a heavy tail. An example of this will be seen in section 4. It therefore seems desirable to use instead

$$T' = \frac{\bar{x}x_{(n)}}{(n - 1)\bar{x}_{GA}^{\ 2}},$$

recognizing that this will differ very little from T when $x_{(1)}$ is small. (The constant $(n - 1)$ is retained only to reduce the dependence of critical values on n.) The statistic T' is proposed as a test statistic—the null hypothesis to be rejected when T' is large—against the ICME alternative, and its use will now be investigated.

4. Use of the Test Statistic

To use the proposed test statistic T', its critical values are needed under the hypothesis that X has an exponential distribution. By design, rejection of the null hypothesis will occur when the distribution is heavier-tailed than the exponential. Since the test is scale-invariant, it can be assumed that $E(X) = 1$ without loss of generality.

The complicated form of $\bar{x}_{GA}$ makes it unlikely that the distribution of T' can be found explicitly. Thus, critical values were estimated for the cases $n = 10$, 15, 20, 25, and 30, by generating 1,000 values of T' for each case on a Hewlett–Packard 9810A programmable calculator. The 10%, 5%, and 1% critical values were then taken as the 90th, 95th, and 99th percentiles of the generated distribution. Because of the randomness in the method, using a sample of size 1,000, we have with 95% confidence:

Quoted level $\alpha = 0.01$ corresponds to actual $0.004 < \alpha < 0.017$
Quoted level $\alpha = 0.05$ corresponds to actual $0.036 < \alpha < 0.064$
Quoted level $\alpha = 0.10$ corresponds to actual $0.08 < \alpha < 0.120$.

With this understanding, critical values are as follows:

	$\alpha = 0.01$	$\alpha = 0.05$	$\alpha = 0.10$
$n = 10$	.40	.369	.362
$n = 15$	.39	.358	.331
$n = 20$	.37	.334	.312
$n = 25$	.36	.322	.298
$n = 30$	.35	.295	.271

Although there is some dependence of critical values on n, the table indicates sufficiently slow variation that interpolation can be used safely.

In order to estimate the power of the test, the same kind of Monte Carlo simulation was performed against the Lomax distributions

$$\bar{F}(x) = \left(\frac{1}{1 + bx}\right)^{1+1/b}$$

for the values $b = 1/7$, 1/5, 1/3, 1, and 3, and for sample sizes $n = 10$ and $n = 20$. Since 1000 runs were again used in the simulation, it should be understood from binomial distribution theory that power values may (at 95% confidence) be off by as much as ±.032, with somewhat better accuracy at very high or very low power. For $n = 10$, estimated power values are as follows:

	$b = 1/7$	$b = 1/5$	$b = 1/3$	$b = 1$	$b = 3$
$\alpha = .01$	.051	.057	.097	.256	.398
$\alpha = .05$	.127	.151	.203	.388	.523
$\alpha = .10$	.149	.181	.234	.407	.553

For $n = 20$, estimated powers are

	$b = 1/7$	$b = 1/5$	$b = 1/3$	$b = 1$	$b = 3$
$\alpha = .01$	.061	.100	.178	.393	.663
$\alpha = .05$	.147	.196	.308	.526	.768
$\alpha = .10$	.229	.281	.388	.627	.819

These power functions are graphed in Figures 1 and 2—as functions of $1/b$ rather than b, for convenience of reading the graph for small values of b while demonstrating the approach toward the exponential distribution at $b = 0$. Thus the abcissa is essentially the degree of the polynomial tail in the density function. It is interesting to note that even for a 7th degree tail, the power is still noticeably above its asymptotic value of α.

(It should be noted that the same data run was used in each case to estimate the power at the three different α levels. Thus, these power estimates are not independent. The estimates for different values of b and N are based on independent simulations, though.)

To illustrate empirical applications, the statistic was applied to three sets of data. The first, from Mielke (1973), is a set of 30 1-day precipitation data collected at Climax, Colorado, by Mielke, Grant, and Chappell (1971):

.085	.030	.060	.080	.110	.100
.495	.260	.085	.220	.390	.045
.010	.020	.035	.055	.190	.015
.025	.065	.125	.010	.125	.020
.015	.010	.005	.070	.100	.030.

These data were found by Mielke to be well fitted by a kappa distribution (v.s.). Applying the proposed test statistic to the data yields

$$T' = .265,$$

and the hypothesis is marginally acceptable ($p \doteq .15$ by extrapolation). Although there is an indication of heavy-tailedness here, it also appears that an exponential-tail hypothesis would not be badly out of line (as was found in the referenced work).

The second set of data consists of 26 observations of precipitation in a Florida meteorological study by Simpson (1972):

129.6	303.8	200.7	978.0	118.3
31.4	119.0	274.7	198.6	255.0
2745.6	4.1	274.7	703.4	115.3
489.1	92.4	7.7	1697.8	242.5
430.0	17.5	1656.0	334.1	32.7
				40.6

With these data,

$$T' = .353,$$

and the null hypothesis can be rejected at the 2% level, approximately (interpolating in the table of critical values). Thus there is somewhat stronger evidence of heavy-tailedness in this case.

Finally, a set of data from Yevjevich (1972) presents annual flows of the Weldon River at Mill Grove, Missouri (1930–59):

108.0	472.0	143.0	441.0	244.0	132.0
53.6	96.5	93.7	386.0	400.0	44.0
585.0	217.0	398.0	567.0	245.0	72.5
98.1	42.7	298.0	122.0	114.0	135.0
40.6	208.0	248.0	151.0	659.0	635.0

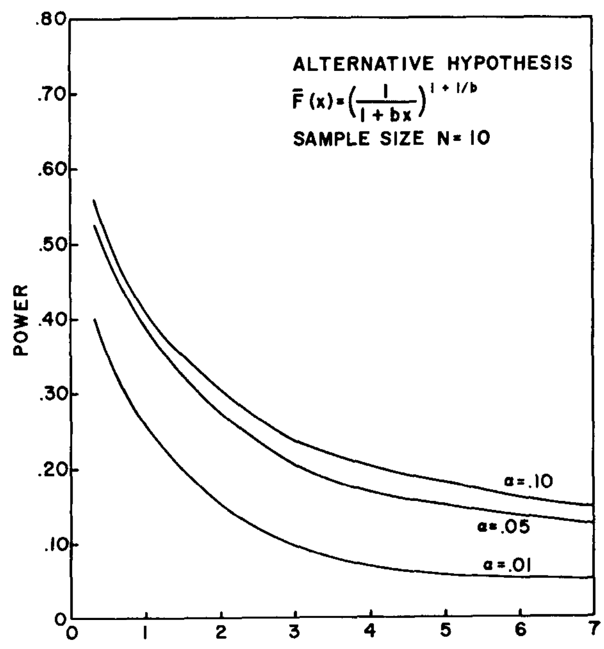

FIGURE 1

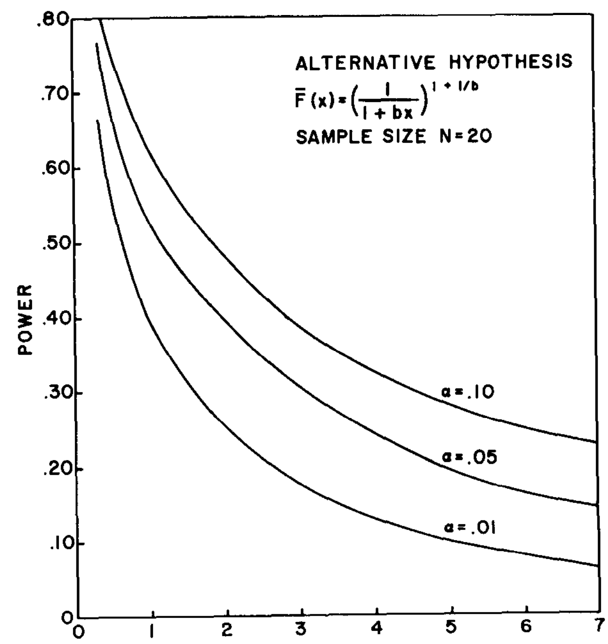

FIGURE 2

For this set of data,

$$T' = .1309$$

and the hypothesis cannot be rejected. This confirms the result of Yevjevich which indicates a gamma distribution. It should be noted, however, that if we had instead used the unmodified statistic T from the previous section, we would have found $T = .365$, and (using results not given here for the distribution of T) the null hypothesis would have been rejected at about the 0.01 level. This phenomenon results from the fact that for these data, $x_{(1)}$ is not small compared to $x_{(n)}/n - 1$; in fact, it is larger. It is this kind of data that argues for the ues of T' rather than T.

5. Graphical Methods

Often a graphical display is the fastest and clearest indication of the form of a distribution. However, for reasons already noted, the usual graphical techniques of a frequency histogram or an empirical *cdf* tend to be uninformative regarding tails of a distribution.

For non-negative data of the type we have been concerned with, a convenient technique is to plot $\log \bar{F}_e(x)$ against x, where $\bar{F}_e(x)$ is the complement of the empirical *cdf*. The technique is convenient, requiring only standard semilog paper; it serves the purpose of amplifying the tail of the distribution; and it often provides an easy visual check for the ICME property.

For the exponential distribution $\bar{F}(x) = \exp(-\lambda x)$, a plot of $\log \bar{F}(x)$ against x will be a straight line emanating downward from the origin. Likewise, the graph of a distribution with an exponential tail, such as the gamma will approach such linearity for large x. ICME distributions, then, will be characterized by graphs that do not approach such linearity and which remain too high. Unfortunately it is difficult to be more precise. This property means that the graph will *tend* to be concave (bending upward) for large x, but it is possible to have temporary interruptions of the concavity without destroying the ICME property.

It is true, and easily verified, that for the ICME distributions the graph of $\log G(x)$ will be concave, where

$$G(x) = \int_x^\infty \bar{F}(t)\,dt.$$

However, $G(X)$ is more often than not difficult to express in closed form or to calculate, and has little intuitive appeal. Thus its use appears less desirable in spite of this mathematical convenience.

The utility of a graph of $\log \bar{F}(x)$ can perhaps be best shown by example. Figure 3 compares the graphs of five distributions:

A: Pareto distribution:

$$f(x) = 2A^2/x^3 \quad \text{for} \quad x > A$$

B: Kappa distribution (see Mielke (1973)):

$$f(x) = A/(A + x)^2$$

C: Exponential distribution:

$$f(x) = A^{-1} \exp(-x/A)$$

D: Gamma distribution:

$$f(x) = A^{-2} x \exp(-x/A)$$

E: One-sided normal distribution:

$$f(x) = 2(\sqrt{2\pi}\,A)^{-1} \exp(-x^2/2A^2).$$

In each case the scale parameter A has been selected to satisfy the condition

$$\bar{F}(10) = .05,$$

so that the five distributions can be compared with respect to their tail behavior around the 95th percentile.

For comparison, Figure 4 shows $\log \bar{F}_e(x)$ for the Florida precipitation data presented in section 4. Despite the small number of observations, a concavity can be seen that is suggestive of the kappa distribution in Figure 3.

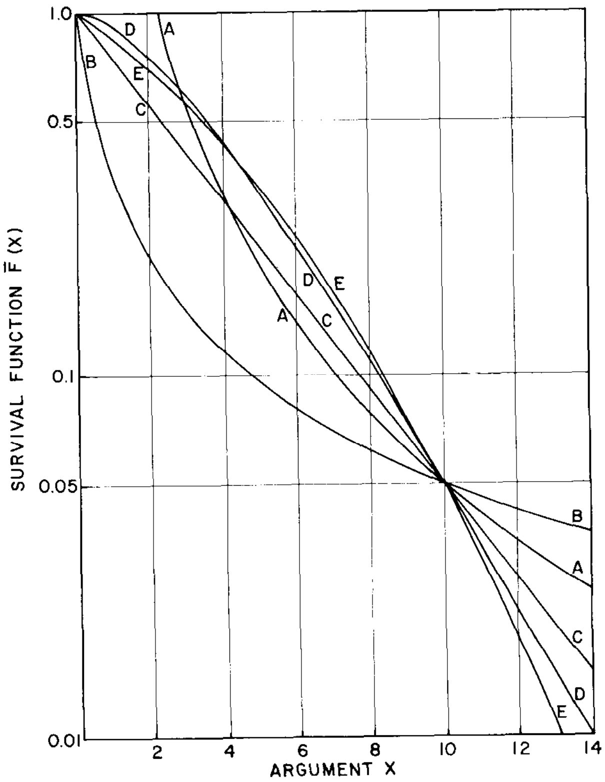

FIGURE 3

6. Alternative Approaches

The analytical and graphical methods discussed here appear to be reasonable ways of testing for ICME, but in the absence of any uniformly most powerful test there are alternatives that should be investigated. One easy test would be based on the statistic

$$T = x_{(n)}/\bar{x},$$

the ratio of the largest observation to the sample mean. It was discussed by Bryson (1968) in connection with tests for *decreasing* mean residual lifetime (or decreasing CME), and in fact provides a most powerful test of the exponential hypothesis against the alternative of a uniform distribution. It is an easy statistic to calculate, is intuitively reasonable, and has a known distribution, being equivalent to a statistic given by Fisher (1950) for testing spectral density significance.

Fercho and Ringer (1972) discuss several tests for exponentiality, of which one appears suitable for the prupose discussed here. This is a statistic proposed by Gnedenko,

$$Q(r) = \sum_{i=1}^{r} (S_i/r) \Big/ \sum_{i=r+1}^{n} (S_i/n - r),$$

where $S_i = (n - i + 1)(x_{(i)} - x_{(i-1)})$ and $x_{(0)} = 0$. This statistic has an F distribution with $2r$ and $2(n - r)$ degrees of freedom if X has an exponential distribution. By choosing $r = n - 1$ or $n - 2$, perhaps, the statistic could be made to depend sufficiently strongly on the tail observations to make it a desirable test statistic against the ICME alternative.

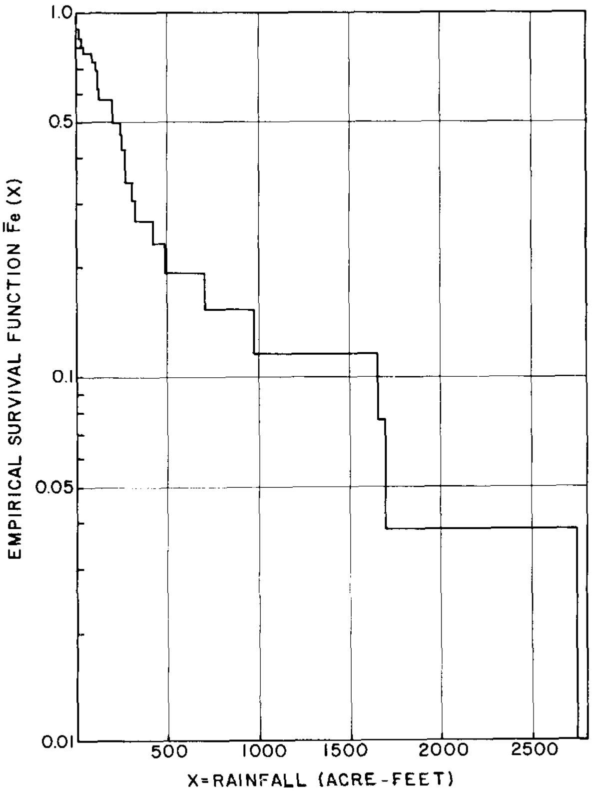

Figure 4

A further subject for investigation should be the behavior of the various tests under exponential-tailed, but non-exponential distributions such as the gamma. Since the gamma distributions (with shape parameter exceeding 1) have decreasing CME, it seems reasonable that a test of the exponential distribution would be a conservative test of gamma distribution; however, this should be verified for any given test.

A final question that could be investigated concerns the existence of the variance in a heavy-tailed distribution. Both the kappa and Pareto families have cases where the variance does and does not exist. As noted in the introduction, certain economic and hydrologic applications seem to call for infinite-variance distributions. However, testing for infinite variance is intrinsically a difficult problem, in that there is no upper bound to finite variances and that sample variances are always finite. There may be some convenient families of distributions that could be investigated from the standpoint of testing for infinite variance, and this problem seems worthy of further study.

References

[1] Barlow, R. E. and Proschan, F. (1965). *Mathematical Theory of Reliability*. John Wiley and Sons, New York.

[2] Boes, D. C. and Salas-La Cruz, J. D. (1973). On the Expected Range and Expected Adjusted Range of Partial Sums of Exchangeable Random Variables. To appear in *J. Appl. Prob.*

[3] Bryson, M. C. (1968). *Topics in Survival Time Distributions*. Ph.D. dissertation, Colorado State University, Fort Collins, Colorado.

[4] Bryson, M. C. and Siddiqui, M. M. (1969). Some Criteria for Aging. *J. Amer. Statist. Assoc. 64*, 1472–1483.

[5] Davis, D. J. (1952). An Analysis of Some Failure Data. *J. Amer. Statist. Assoc.* 47, 113–150.

[6] Epstein, B. and Sobel, M. (1953). Life Testing. *J. Amer. Statist. Assoc.* 48, 485–502.

[7] Fercho, W. W. and Ringer, L. J. (1972). Small Sample Power of Some Tests of Constant Failure Rates. *Technometrics 14*, 713–724.

[8] Fisher, R. A. (1950). *Contributions to Mathematical Statistics*, Paper XVI. John Wiley and Sons, New York.

[9] Granger, C. W. J. and Orr, D. (1972). "Infinite Variance" and Research Strategy in Time Series Analysis. *J. Amer. Statist. Assoc. 67*, 275–285.

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

[11] Hollander, M. and Proschan, F. (1972). Testing Whether New is Better Than Used. *Ann. Math. Statist. 43*, 1136–1146.

[12] Johnson, N. L. and Kotz, S. (1970). *Continuous Univariate Distributions*. Houghton-Mifflin Company, Boston.

[13] MALIK, H. J. (1970). Estimation of the Parameters of the Pareto Distribution. *Metrika 15*, 126–132.

[14] MIELKE, P. W. (1973). Another Family of Distributions for Describing and Analyzing Precipitation Data. To appear in *J. Appl. Meteorology*.

[15] MIELKE, P. W., GRANT, L. O., and CHAPPELL, C. F. (1971). An Independent Replication of the Climax Wintertime Orographic Cloud Seeding Experiment. *J. Appl. Meteorology* 10, 1198–1212.

[16] SIMPSON, J. (1972). Use of the Gamma Distribution in Single-Cloud Rainfall Analysis. *Monthly Weather Review* 100, 309–312.

[17] WAGNER, H. M. (1969). *Principles of Operations Research*. Prentice-Hall, Inc., Englewood Cliffs, New Jersey.

[18] YEVJEVICH, V. (1972). *Probability and Statistics in Hydrology*. Water Resources Publications, Fort Collins, Colorado.