26 From Diversity to Volatility: Probability of Daily Precipitation Extremes

Anna K. Panorska[1], Alexander Gershunov[2] and Tomasz J. Kozubowski[1]

[1] Department of Mathematics and Statistics, University of Nevada Reno
[2]Climate Research Division, Scripps Institution of Oceanography

Abstract. A sensible stochastic model is required to correctly estimate the risk associated with daily precipitation extremes. The same requirement holds for studying high-frequency precipitation extremes in the context of climate variability and change. Results derived from probability theory were used to develop an efficient automated scheme to distinguish between heavy and exponential precipitation probability density function (PDF) tails in hundreds of daily station records spanning five decades over the North American continent. These results suggest that, at a vast majority of the stations, daily extreme precipitation probabilities do not decay exponentially, but more closely follow a power law. This means that statistical distributions traditionally used to model daily rainfall (e.g. exponential, Weibull, Gamma, lognormal) generally underestimate the probabilities of extremes. The degree of this distortion, i.e. volatility, depends on regional and seasonal climatic peculiarities. By examining geographical and seasonal patterns in extreme precipitation behavior, the authors show that the degree of volatility is determined regionally by the diversity in precipitation-producing mechanisms, or storm type diversity. Exponential tails are geographically limited to regions where precipitation falls almost exclusively from similar meteorological systems and where light probability tails are observed in all seasons. Topography plays an important role in flattening or fattening PDF tails by limiting the spatial extent of certain systems while orographically altering their precipitation amounts. Results presented here represent the first logical step towards choosing appropriate PDFs at various locations by specifying their regionally relevant family. Heavy tailed models are generally superior to those from the exponential family and can lead to more realistic estimates of extreme event probabilities, return periods, n-year events, and design limits. The correct choice of PDF is essential to safe engineering design, hazard assessment and other applications, as well as for fostering further investigations of hydrologic weather extremes and climate.

1 Introduction

High frequency hydrologic extremes are notoriously volatile and, therefore, difficult to model with traditional probability density functions (PDFs), those with exponentially decreasing probabilities of extreme occurrences. Nevertheless, because they are traditional, exponentially tailed PDFs are typically used to model hydrologic data with the result that observed extremes are left out as improbable events or outliers. This situation leads to traditionally inadequate treatment of daily precipitation extremes in climate research as well as to underestimation of extreme event probabilities in applications which require estimates of environmental risk. Attempts to determine the correct PDF for daily rainfall amounts have so far been either locally focused or assumed simple models for rainfall production universally with the confusing consequence that we still do not know what environmental factors determine the structure of probability tails for daily precipitation. Contradictory results using inconsistent methodologies applied at various locations yield various precipitation PDFs with either exponential or heavy (power law) tails. In this chapter, we try to solve the classification problem in a mathematically, methodologically and spatially consistent way by letting daily precipitation data at hundreds of stations across North America show us the nature of its tails.

The correct choice of PDF is not trivial. For rainfall, a hodgepodge of available theoretical PDFs with either exponential or heavy tails has so far complicated this choice. The data-based approaches to choosing correct PDFs for daily rainfall mostly utilized so far involved fitting specific PDFs to data at limited locations (e.g. Katz et al. 2002). Because, such data-based approaches are methodologically inefficient for choosing correct PDFs at a multitude of available observing stations with daily rainfall, geographic variability in extreme rainfall behavior has not been examined. Alternatively, attempts to determine a "universal" PDF for modeling daily rainfall everywhere have also been made. Most recently, Wilson and Toumi (2005) assumed a simplified model for rainfall production without regard to storm type and then mathematically derived a stretched (i.e. double or extended) exponential PDF as the universal model for extreme rainfall events. In this type of work, data is used as a means to validate the simplified theoretical result, but not to achieve it. Taken in their entirety, previous results do not agree on the existence of a universal PDF for daily rainfall. Nor, do previous results agree on which family of PDFs (i.e. exponentially- or heavy-tailed) best describes rainfall extremes at most locations. We do not know the spatial patterns of applicability of different PDFs for extreme rainfall modeling and, needless to say, we do not have an understanding of the physical processes that may be responsible for any such possible spatial variability or dominance.

In this chapter, we develop and apply an efficient methodology to determine whether exponentially- or heavy-tailed *families* of PDFs best describe the data at hundreds of stations across North America. This data-based *family* classification problem is, in our thinking, the first step towards understanding hydrologic weather extremes in a climatic context. Fitting specific PDFs at specific locations should be greatly facilitated once such classification is done. Moreover, even a simple visual analysis of the classification result applied to hundreds of stations across the North American continent, should illuminate the physical causes for any possible geographic patterns in the dominance of one family over another.

Statisticians have tried to fit heavy tailed models to hydrologic data. Heavy tailed behavior has been documented in precipitation and streamflow measured at specific locations or regions for time series of extreme values, e.g. annual peak record streamflow (Smith 1989, Katz et al. 2002) and precipitation (Katz et al. 2002), as well as threshold exceedances for streamflow (Anderson and Meerschaert 1998). Heavy tails have also been *suggested* as reasonable models for daily precipitation amount at a large number of U.S. stations (Smith, 2001). However, a formal and efficient classification scheme for heavy vs. exponential probability tails has not been developed and, in climate and hydrology research, exponentially tailed PDFs are still typically used to model daily rainfall, including extremes (e.g. Tsonis 1996, Groisman et al. 1999a, Zolina et al. 2004, Wilson and Toumi 2005).

Heavy tailed distributions such as Pareto (Johnson et al. 1994) or the stable laws (Samorodnitsky and Taqqu 1994) arise naturally as approximations to sums of random quantities such as precipitation accumulated over time. Pareto distribution, via the Peak Over Threshold theory, is also a natural model for excesses over high thresholds. Formally, a random variable X is heavy tailed if $P(X > x) = cx^{-\alpha}$, where c and α are positive constants, for large values of x. Alternatively, we can say that heavy tailed distributions have "power" tails. The most important, essentially qualitative, difference between the heavy and exponentially tailed random quantities is in the size of their large percentiles relative to the majority of the data. The heavy tailed distributions have much larger high percentiles than the exponentially tailed laws, which reflects their higher volatility. We aim to quantify volatility of daily precipitation observed at hundreds of stations over North America and to examine the spatial structure of the result. To quantify volatility and to classify station precipitation into exponential or heavy tailed families of distributions, we use a novel approach stemming from a combination of physical reasoning and a careful examination of the connection between data and mathematics.

We approach the heavy vs. exponential tail classification problem in the context of peaks over threshold (POT) methodology and derive a new statistical test based on the theory of maximum likelihood ratios. Theory, presented in Section 2, implies a broader impact for our test in that it classifies entire PDFs of daily precipitation into exponentially- and heavy-tailed families of distributions. In section 3, we find that daily precipitation probability tails are heavy, not exponential, at a great majority of stations. Moreover, a geographic dependence of volatility is discovered, whereby the heaviest tails occur where precipitation is produced by a variety of storm types, while exponential or light tails occur where diversity in precipitation-producing systems is limited. Important orographic and seasonal controls on precipitation volatility emerge. A practical impact of these results on estimation of extreme event return periods is briefly considered. Section 4 presents a summary of results and discusses some broad implications for applications and climate research.

2 Methodology and Data

We want the data to show us the nature of its tails: are they exponential or heavy? Aiming to do such a comparison, we are immediately confronted by the fact that many choices exist for the specific form of the PDF with exponential and heavy tails. We consider a random variable X to be heavy tailed if $P(X > x) = cx^{-\alpha}$, where c and α are positive constants, for large values of x. Note, that according to this definition, none of the distributions traditionally used to model precipitation (e.g. Gamma, lognormal, Weibull, double, stretched or extended exponential, etc.) are heavy tailed. Examples of heavy tailed distributions include Pareto, Cauchy, Fréchet, and stable laws. To resolve the problem of distributional multiplicity, i.e. to avoid having to choose specific PDFs from the exponential and heavy-tailed families, we applied the results of the peak-over-threshold (POT) theory. Since our primary interest is in estimating probabilities of large (extreme) events, the POT method helps focus the search for the reasonable models and provides a rigorous mathematical foundation for the results.

2.1 Mathematical theory and statistical test

The POT method involves examination of the data falling above a threshold. For any data value X and threshold u, $X^{[u]}$ is the *exceedance*, that is the (conditional) value of X given that X exceeds u. We consider *excesses* over the threshold u, i.e. $X^{[u]}$-u. The Balkema – de Haan – Pickands theorem (Balkema and de Haan 1974, Pickands 1975) provides the limiting distribution of excesses. The theorem states that when the threshold (u) increases, the distribution of the excesses ($X^{[u]}$-u) converges to a Generalized Pareto (GP) distribution. Any GP distribution has to be one of the following three kinds: exponential, Pareto or Beta. The importance of this result for practical applications is: *no matter what the specific original distribution of X is, the excess, $X^{[u]}$-u, over threshold u has (approximately) one of only three distributions.* The three distributions of the excesses correspond to the tails of the original distribution of X. If X has exponential tail, then $X^{[u]}$-u will have an approximately exponential distribution. If X has heavy tail, then $X^{[u]}$-u will have an approximately Pareto distribution. Beta distribution has finite support and we do not consider it here. Our problem of finding the correct PDF for describing daily extreme event probabilities reduces, therefore, to classifying excesses at a particular station. We seek a decision rule, ideally a formal statistical test, for classifying the precipitation excesses over threshold into either exponential or Pareto models.

We approached this problem using ideas from the theory of likelihood ratio tests (Lehmann, 1997). Formally, we test the null hypothesis

H_0: data comes from an exponential distribution,
versus the alternative
H_1: data comes from a Pareto distribution.

The approach is to consider the ratio of the maxima of the likelihoods of the observed sample under the Pareto or exponential (in the numerator) and exponential (in the denominator) models. The logarithm[1] of the likelihood ratio, the L statistic, is:

$$L = \log\left(\frac{\max(\sup_{\alpha>0,s>0} L_{Pareto}(\vec{x}|\alpha,s), \sup_{\sigma>0} L_{\exp}(\vec{x}\mid\sigma))}{\sup_{\sigma>0} L_{\exp}(\vec{x}|\sigma)} \right),$$

where $\vec{x}$ is the observed sample of excesses and $L_{Pareto}(\vec{x}|\alpha,s)$ and $L_{\exp}(\vec{x}|\sigma)$ are the likelihood functions of the sample under Pareto and exponential models, respectively. We use a Pareto distribution with the survival function S(x) = P(X>x) = $(1/(1+x/s\alpha))^{\alpha}$ and exponential distribution with the survival function S(x) = P(X>x) = $\exp(-x/\sigma)$. To compute L, both likelihoods are maximized first (via maximum likelihood estimates, MLEs, of the parameters), and then the natural logarithm of their ratio is taken as the likelihood ratio statistic.

To help the reader implement the numerical routines necessary for the computation of L, we include several formulas required for the maximum likelihood procedure. Let $\vec{x}=x_1, \ldots, x_n$ be a sample of size n. The supremum of the natural logarithm of the exponential likelihood for this sample is

$$\sup_{\sigma>0} \log(L_{\exp}(\vec{x}\mid\sigma)) = n(-\log(\bar{x})-1),$$

where $\bar{x}$ is the sample mean. The natural logarithm of the supremum of the Pareto likelihood is

$$\sup_{\alpha>0,s>0} \log(L_{Pareto}(\vec{x}\mid\alpha,s)) = n(\log(\hat{\alpha})-\log(\hat{s})-1-1/\hat{\alpha}),$$

[1] All logarithms (log) are natural logarithms.

where $\hat{\alpha}$ and $\hat{s}$ are the MLEs of α and s for the Pareto likelihood. The computation of $\hat{s}$ requires numerical maximization of the function

$$u(t) = -\log\left(\frac{\sum_{i=1}^{n}\log(1+x_i t)}{n}\right) + \log(t) - \frac{1}{n}\sum_{i=1}^{n}\log(1+x_i t),$$

with respect to t > 0. If $\hat{t}$ is the maximum of u(t), then the MLE of s is $\hat{s}=1/\hat{t}$. The MLE of α is

$$\hat{\alpha} = \frac{1}{(1/n)\sum_{i=1}^{n}\log(1+\frac{x_i}{\hat{s}})}.$$

The properties of the test, proofs and more details on the optimization process will be published separately (Kozubowski et al., manuscript in preparation).

The limiting distribution of the test statistic L is a mixed distribution, with an atom at zero and distribution given via the distribution of Y=2L:

$$Y = \begin{cases} 0 & with \quad probability \quad 1/2 \\ \chi_1^2 & with \quad probability \quad 1/2 \end{cases}$$

Using the formula above, we can compute (1-γ)100 percentiles for Y as follows. Any percentile below the 50th percentile is 0. For 0 ≤ γ ≤ 0.5, c = (1-γ)100 percentile for Y if P(Y ≤ c) = 1-γ, which is equivalent to P(χ_1^2≤ c) = 2(1-γ)-1 and thus c can be found in the χ_1^2 tables. Then, the (1-γ)100 percentile of L is c/2.

The (1-γ)100 percentiles of L provide the critical numbers for our test on the significance level γ. The test is one-sided and we reject the null hypothesis if the computed value of the test statistic exceeds the critical number. We have computed some common percentiles for the distribution of L under the null hypothesis for different sample sizes and for the limiting case. The percentiles for finite sample sizes were computed via Monte Carlo simulation with 10,000 samples of a given size from the exponential distribution (Table 26.1).

Table 26.1. The entries are the (1-γ)100 percentiles of the distribution of L for various sample sizes and the limiting distribution (last row) of L. These are also critical numbers for testing our hypothesis on different significance levels γ.

CRITICAL VALUES OF L

Sample size	Significance level, γ [and confidence (1 - γ)*100]			
	0.01 [99%]	**0.02** [98%]	**0.05** [95%]	**0.1** [90%]
10	1.71128	1.1561	0.62701	0.25703
50	2.15057	1.5768	0.89045	0.48852
100	2.23171	1.71583	0.94963	0.55706
500	2.45615	1.85952	1.18044	0.70439
1,000	2.51298	1.92766	1.22475	0.71376
5,000	2.62019	1.97122	1.27738	0.76095
10,000	2.70307	2.0285	1.30714	0.80146
∞	2.70595	2.10895	1.35275	0.82120

2.2 Intuitive insight

On a more intuitive note, in the Pareto case, the $\boldsymbol{\alpha}$ parameter determines the thickness of its tail and is of primary importance. As $\boldsymbol{\alpha}$ decreases to zero, the tail of the Pareto distribution becomes heavier (larger volatility) causing the probabilities of extremes to increase. The scale parameter $\boldsymbol{s}$ is of secondary importance. In the exponential case, $\boldsymbol{\sigma}$ is the scale parameter. Statistic L is scale invariant. Note, that as the tail parameter $\boldsymbol{\alpha}$ increases, the Pareto survival function converges to the exponential one. We included the limiting case, $\boldsymbol{\alpha} = \infty$, in the numerator of L. This convergence of Pareto to exponential adds particular difficulty to differentiating between the two distributions as $\boldsymbol{\alpha}$ increases (Fig. 26.1).

The intuition behind this approach is as follows. If a sample (i.e. excesses over threshold) comes from a Pareto distribution, statistic L is expected to be positive because the maximum Pareto likelihood is likely to be larger than the exponential one. Thus, the ratio will be greater than one (L > 0). If the sample comes from the exponential distribution, then the Pareto likelihood function is likely to be maximized by the limiting value of $\boldsymbol{\alpha}$ (i.e. $\boldsymbol{\alpha}=\infty$), corresponding to the exponential distribution so that the ratio of the two likelihoods is 1 and its natural logarithm is 0. The distribution of L for Pareto samples with various values of $\boldsymbol{\alpha}$ and for exponential samples is shown in Fig. 26.1.

While, in theory, L should be zero for every exponential sample and positive for any Pareto sample, L has its own variability and the theoretical results do not hold exactly in practice. The main question when considering a statistic for a classification problem is: Does the statistic differentiate between the two distributions in question. We assessed the viability of L using simulations because the true theoretical distribution of L is not known. The results presented in the form of boxplots in Fig. 26.1 confirm that L differentiates between the

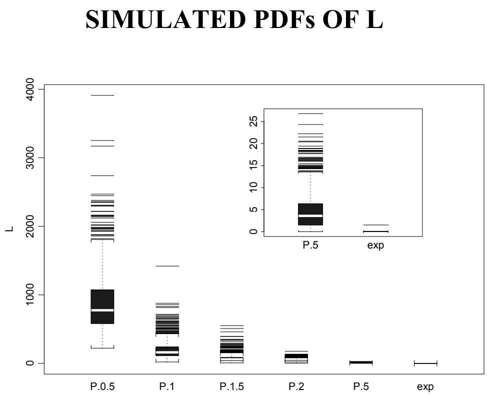

Fig. 26.1. The graph contains 6 boxplots of simulated distributions of L. The first five boxplots were done using 10,000 observations of L from Pareto samples of size 1,000 with α varying from 0.5 (first boxplot) to 5 (second to the last boxplot). The last boxplot corresponds to 10,000 observations of L from exponential samples of size 1,000. The inset blows up the last two boxplots.

Note that the distribution of L for Pareto samples with small $\boldsymbol{\alpha}$ is very different than that for the exponential samples. This provides intuitive evidence that L is able to differentiate between exponential and Pareto distributions where they differ considerably (small $\boldsymbol{\alpha}$). When the Pareto and exponential distributions become very close to each other (large $\boldsymbol{\alpha}$), it follows that any probabilities computed from the two (close) models will be very close, so the question of which distribution is better to use becomes mute.

The convergence of Pareto to exponential as $\boldsymbol{\alpha}$ increases is observed as the convergence of the distribution of L for Pareto samples with increasing α to the distribution of L for the exponential samples (Fig. 26.1). Additionally, as $\boldsymbol{\alpha}$ decreases the L statistic becomes "larger on average". This can be visualized by looking at the boxplots and at the values of the median of L that increase as $\boldsymbol{\alpha}$ decreases. This (average) monotonicity of L leads to the conclusion that larger values of L provide, on average, stronger evidence for Pareto distribution. Smaller $\boldsymbol{\alpha}$'s (larger L values, on average) imply heavier tails, which imply larger probabilities of extremes.

We programmed and tested the computation of L from the numerical viewpoint. We also ran several simulation studies and applied the method to the precipitation data with various thresholds. Although threshold choice remains largely intuitive so far, we obtained fairly robust and meaningful results for daily precipitation with thresholds ranging from the 50th to the 95th percentiles. Below, we present results based on the 75th percentile threshold, which allows a reasonable amount of data for statistical testing in most regions and seasons. In discussing results below, we focus primarily on features generally robust to threshold choice. We do note, however, that optimal threshold choice is recognized as a difficult and still open problem (e.g. Smith 1987, Davidson and Smith 1990, Smith 1994, Gross et al. 1994).

2.3 Data

We have computed L for a subset of over ten thousand available station records of daily rainfall distributed across Canada (Vincent and Gullet 1999), U.S. (Groisman et al. 2004) and Mexico (Miranda 2003) all quality controlled and homogenized at the National Climatic Data Center (NCDC). Five hundred sixty best quality stations were selected to give a reasonable coverage (at least one station within a 70km radius) over North America. The selection used the most stringent quality standards in well-sampled regions and relaxed these standards to include all stations in regions with the poorest coverage (i.e. Alaska, Northern Canada, and parts of Mexico). In mountainous regions where coverage allowed, we always included the highest elevation station along with the best quality station within the 70km radius. At least 80% of the data was required to be present at all stations for the common observational time period: January 1, 1950 to December 31, 2001. In what follows, we compare fits to excesses over the local 75th percentile derived from Pareto to those obtained from the widely used exponential. Figure 26.2 shows the log likelihood ratio statistic (L) at each station.

3 Results

Figure 26.2 is color coded according to the heaviness of the local precipitation tail. This result suggests a spatially coherent dependence of tail type on geographical location with respect to climate. Figure 26.2a presents the absolute magnitude of L. Figure 26.2b more clearly presents the classification result in terms of the statistical confidence with which the null hypothesis (H_o) of exponential tails can be rejected in favor of the heavy-tailed alternative (H_1). The latter result takes sample size into account by interpolating the critical values presented in Table 26.1.

First of all, we see that the vast majority of local tails are strictly speaking non-exponential. In fact, at 81% of the stations, H_o can be rejected with 95% confidence. The degree of departure from the exponential model depends on the peculiarities of regional climatic regimes. Precipitation at stations with a large variety of meteoro-logical influences (i.e. frontal, thunderstorm, tropical cyclone, intense and organized convection) tends to display highly non-exponential tails. This occurs especially where extreme weather interacts with noteworthy topography (e.g. the Mexican Gulf coast, the front range of the Rockies), the regions of hurricane landfalls (e.g. Gulf and East Coasts of the U.S.), and along Tornado Alley. In the Great Plains and the Midwest, especially the north, heavy precipitation events can appear more extreme relative to typically light winter snow accumulations.

LOG LIKELIHOOD RATIO TEST STATISTIC (L), ALL DATA

a) L values, absolute magnitude **b)** L values, H_o rejection certainty in %

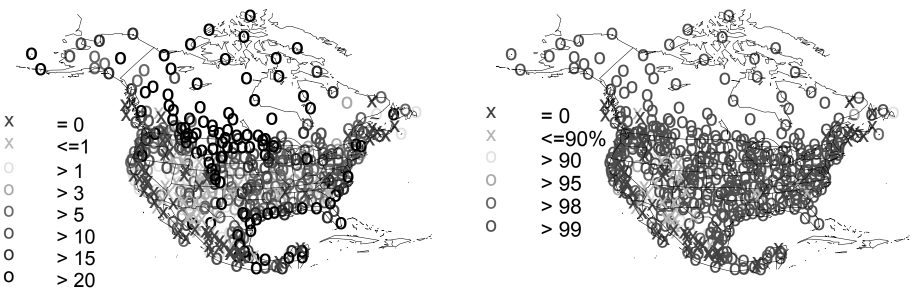

Fig. 26.2. Log likelihood ratio (L) computed for daily excesses over local 75th percentile at each of the 560 stations. **(a)** Values close to zero, L <= 1 (blue and green x's) represent approximately exponential tails, while yellow, red and black circles represent progressively heavier tails. **(b)** Confidence level, $(1 - \gamma)*100$, for rejecting the null hypothesis (H_o) of exponential tails. Blue x's represent exponential tails, green x's represent stations at which H_o cannot be rejected with reasonable (90%) confidence. Yellow and progressively redder circles represent stations at which H_o can be rejected with 90, 95, 98 and 99% confidence in favor of the Pareto alternative. H_o can be rejected at 81% of stations with 95% confidence. (A color version of this figure appears between pages 196 and 197).

Exponential behavior occurs where precipitation is reasonably frequent and preferably of one type, i.e. generated by similar type systems. For example, the exclusively convective summer precipitation along the central Mexican plateau as well as isolated regions of almost entirely frontal-system-generated precipitation along the windward side of gently sloping topography (e.g. Sierra Nevada and Rocky Mountains) exhibit exponential tails. Topography appears to play an important role. Notably, precipitation volatility increases abruptly eastward of the Continental Divide at the steep Front Range of the Rockies.

A simple way to categorize precipitation by frontal and convective regimes is to consider winter and summer data separately. Transitional seasons' daily accumulations are interesting as well, as they present preferential mixtures of storm types. Figure 26.3 presents L computed for the December – February (DJF), March – May (MAM), June – August (JJA), and September – October (SON) seasons separately.

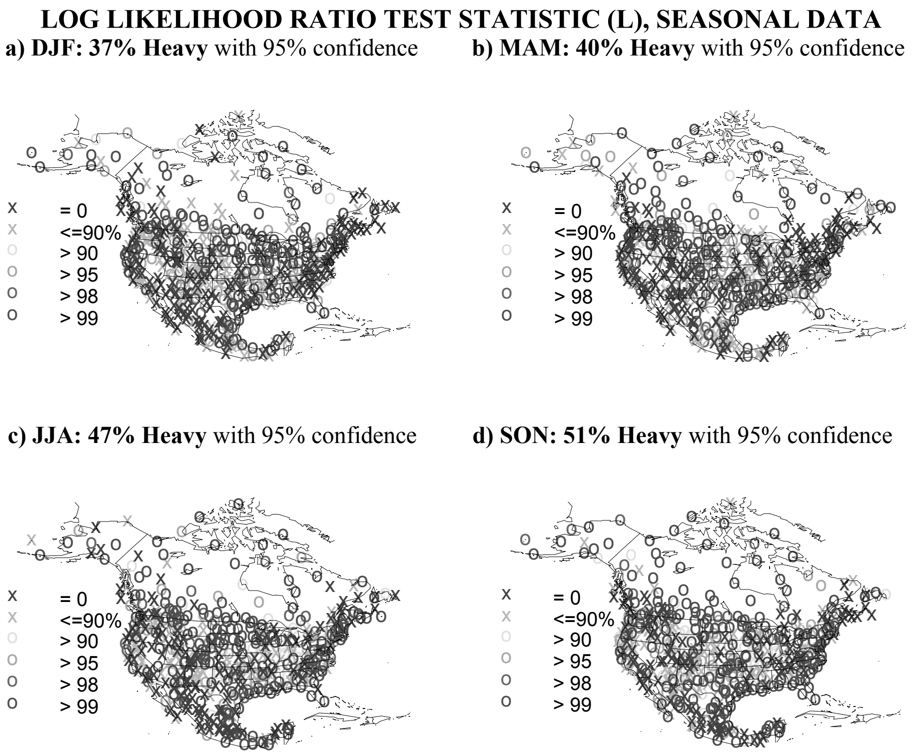

Fig. 26.3. Log likelihood ratio (L) test statistic, as in Fig. 26.2b, but for seasonal data. Plates (a) through (d) show H_o rejection confidence for winter, spring, summer and fall, respectively. Percentage of all stations where exponential tails can be rejected in favor of Pareto with 95% confidence is given in the title for each plate. (A color version of this figure appears between pages 196 and 197).

Precipitation generated by frontal-systems over the mountainous West and the East Coast in winter (Fig. 26.3a) exhibits exponential tails, while the mostly light and moderate cold snow accumulations mixed in with occasional warmer, wetter and heavier events display mostly heavy tails in the northern plains and the Midwestern US. The high Mexican plateau, where it does not rain much in winter, shows mostly exponential tails, while the Mexican part of the Gulf Coast (from southern Texas south to the Isthmus of Tehuantepec), where *nortes* or “northers” bring much of the heaviest winter precipitation exhibits heavy tails. Spring (Fig. 26.3b) mixes convective rainfall with frontal precipitation and brings heavy tails along the US

Gulf Coast and the Deep South. The Continental Divide, visible in all seasons, is best delineated with a strong volatility gradient in spring. Snow accumulations from eastward-moving midlatitude cyclones diminish abruptly east of the Continental Divide, while heavy precipitation occurs here when moist Gulf air is entrapped into the leading edge of south-moving surface anticyclones and slammed against the Front Range.

In summer and fall (Fig. 26.3c,d), heavy tails outline the entire Gulf and East Coasts, where convective rainfall is mixed with that generated by land-falling tropical storms (plus frontal precipitation in fall). The southwestern monsoon fattens precipitation tails in northwestern Mexico and Arizona in summer, while occasional wayward eastern Pacific tropical storms extend heavy tails into Baja California Sur, neighboring Sinaloa, coastal Sonora, and possibly as far north as southern California State in the fall. Along the west coast, the heaviest tails occur in the fall, where typically moderate precipitation amounts can turn very heavy in the occasional early wet-season (most recently, fall 2004 was one of those seasons). The Appalachian range is most clearly outlined with a volatility gradient in the fall. The northwest slope is protected from tropical storm intrusions during the hurricane season, while these storms occasionally reinforce typically lighter precipitation from the more common local and frontal systems southeastward of the range. In general, orography can produce volatility gradients by confining heavy precipitation from specific storm systems to one side of the mountain range, while orographically enhancing their precipitation efficiency. Orography affects storm-type diversity and this effect is typically seasonal.

Generally, L computed for all precipitation data (Fig. 26.2) is an aggregation of seasonal values, reflecting the volatility of precipitation in different seasons (Fig. 26.3). All seasons' precipitation tails are more likely to be heavy at any given station and are heavier than those observed in any one season. For all-season precipitation (Fig. 26.2a), the heaviest tails occur in regions where tails tend to classify heavy in all seasons. These regions are the Gulf Coast, Midwest and the northern plains, and the steepest part of the Front Range of the Rocky Mountains. Exponential tails exist only in regions where precipitation tails are light in all seasons: the gentle west facing slopes of the high western mountains where precipitation is mostly frontal and orographically enhanced, the high plateau of central Mexico where rainstorms are almost exclusively local convective, and at the Canadian Atlantic and Pacific coasts, around drizzly Nova Scotia and Queen Charlotte Islands. Most other regions display overall heavy tails (Fig. 26.2b). The reader is invited to examine Figs. 26.2 and 26.3 to verify that overall heavy tails (Fig. 26.2b) occur in regions where heavy tails appear in at least one of the four seasons. This observation is absolutely consistent with theory, according to which, heavy tails can result by mixing distributions, as long as at least one of those distributions is heavy-tailed.

To help get a visual, intuitive feel for what the measure L represents with respect to the form of the probability tails and their fit to the data, Fig. 26.4 shows probability plots for stations selected to represent specific precipitation regimes and to roughly span the range of L. These probability plots show how heavy and extreme events are accounted for by exponential and Pareto models. Figure 26.4 presents graphical pairs, scatter plots of the quantiles of the empirical distribution against the quantiles of a theoretical model: exponential or Pareto. As in the calculation of L (Figs. 26.2 and 26.3) the daily station data used here are excesses over the local 75^{th} percentile of the daily amounts recorded on wet days (i.e. heavy to extreme events). In this type of plot, the best model is the one that scatters along a straight line.

PROBABILITY PLOTS AND L VALUES FOR SELECTED STATIONS

a) Sacramento, L = 1.60

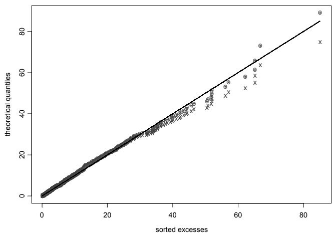

b) Nashville, L = 3.15

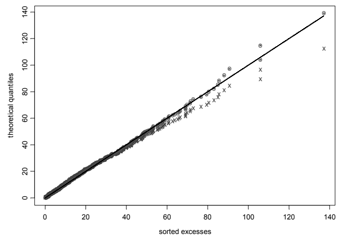

c) St. Louis, L = 4.93

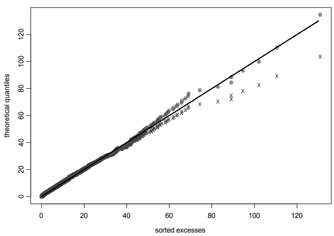

d) Houston Hobby Airport, L = 15.2

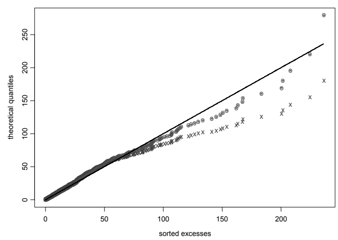

e) Fargo WSO AP, L=28.6

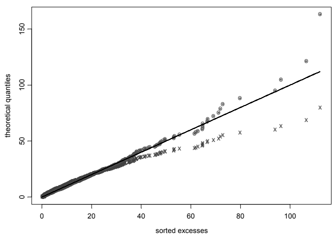

f) Miami WSCMO Airport, L = 41.8

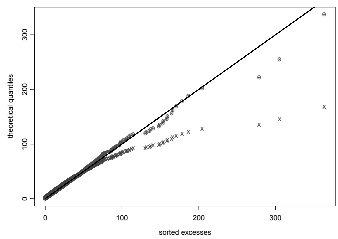

Fig. 26.4. Probability plots for excesses over threshold (75^{th} percentile, see Table 26.1) at selected stations arranged in order of increasing L (see titles for individual plates). Sorted observed excesses displayed in mm along the x-axis are plotted against the corresponding theoretical quantiles derived from the fitted exponential (blue x's) and Pareto (red o's) models.

Table 26.2 presents precipitation statistics at these selected stations for the common observational period 1950–2001: L, probability of precipitation, the 75th percentile of the daily climatology on days with precipitation, maximum recorded daily total, the estimated 100-year event assuming exponential and Pareto tails, and the Pareto probability of exceeding the 100-year event estimated assuming an exponential tail. The thrust of this chapter is on classification. The estimation results presented in the last two columns of Table 26.2, although not rigorous, are shown to illustrate the approximate magnitude that classification can have on natural hazard estimation.

As a rule, where L ≈ 0, both exponential and Pareto tails fit the excesses equally well. We start with an example of a modestly-tailed station, Sacramento FAA AP (40.52°N, 123.82°W, 640m, Fig. 26.4a), located in the Central Valley of California, where rain is produced almost exclusively by frontal systems in the winter half year. Both exponential and Pareto models fit these data reasonably well and L is close to zero. However, the shape of the Pareto tail still differs appreciably from the exponential, especially beyond the observed range. In the 52 years of analyzed data at Sacramento, the 75th percentile of daily total precipitation was 10.7 mm and the maximum was 96 mm. Although both models seem to fit the observed data well, model choice can lead to significant differences in inferences regarding the probabilities of hitherto unobserved extremes. For example, based on the observed

Table 26.2. Precipitation statistics at selected stations for the common observational period 1950 – 2001: L; probability of precipitation (i.e. % of days with recorded precipitation); 75th percentile of daily total on days with precipitation; maximum recorded daily total; the estimated 100-year event assuming exponential and Pareto tails; and the Pareto probability of exceeding the exponential 100-yr event. The last column can be interpreted as the factor by which the 100-yr event estimated assuming exponential tail is more likely to occur assuming Pareto tail. Alternatively, the Pareto return period for an exponential 100-yr event is 100 years divided by the value in the last column at a specific station.

PRECIPITATION STATISTICS AT SELECTED STATIONS

Station	**Log likelihood ratio (L)**	**P[p>0] (%)**	**75th%-ile ($p_{p>0}$) (mm)**	**Max$_{obs}$ (p_{max}) (mm)**	**100yr event** Exp & Pareto **(mm)**	**Pareto P[$p>p_{exp}^{100}$] (%)**
Sacramento	1.60	16	10.7	96	85 & 99	2.3
Nashville	3.15	26	16	153	127 & 154	3.4
St. Louis	4.93	30	11.2	142	114 & 144	4.1
Houston	15.2	27	16.3	253	195 & 292	6.5
Fargo	28.6	27	5.8	118	85 & 167	12.0
Miami	41.8	36	13.7	377	181 & 346	9.8

record and assuming the exponential (Pareto) model, the estimated 100-year event is 85 (99) mm/day. The Pareto model implies that the 100-year event estimated by the exponential model is 2.3 times as likely to occur, i.e. a 43-year event.

Exponential tails become progressively less adequate as we consider stations with progressively larger L. This is readily visible on the probability plots (Fig. 26.4b-f) as well as by comparing expected values of 100-year events with the observed 52-year station maxima (Table 26.2). Pareto tails perform reasonably well, certainly better by comparison with the exponential for all stations considered. This is not to say that they perform perfectly well, though. At some stations, Pareto tails tend to overestimate the observed probabilities of extremes, as exemplified by extreme values recorded at Fargo, North Dakota (Fig. 26.4e). At others, they may somewhat underestimate probabilities of the largest extremes, e.g. Miami (Fig. 26.4f). It is clear, however, that the heavy tail model, as a rule, outperforms the exponential.

Heavy and extreme daily precipitation at a vast majority of the stations is more in line with the Pareto model. The comparison of 100-year events estimated using the exponential and Pareto models with the observed 52-year station maxima (Table 26.2) reinforces the choice of heavy over exponentially-tailed models to describe probabilities of extreme precipitation events. Assuming, as our results suggest, that the Pareto model is correct in most cases, the choice of an exponentially-tailed PDF may lead to a severe underestimation of extreme precipitation risk. In Miami, for example, the 100-year event estimated by the exponential model corresponds to a 10.2-year event as estimated by Pareto, i.e. is 9.8 times more likely to occur.

4 Summary and Conclusions

The statistical methodology adapted here allows us to simply evaluate whether the tails of an empirical distribution are exponential or heavy, regardless of the specific form of the PDF from the exponential or heavy tailed families of distributions. So, it is not necessary to test whether the data is or is not exponential, Gamma, Gumbel, Weibull, Double (or sometimes “stretched” or “extended”) exponential or any other exponentially-tailed PDF commonly used to model daily rainfall. These are accepted or rejected under the broad umbrella of the exponentially tailed family of distributions. Although theoretical questions related to the choice of threshold remain to be solved, the Pareto/exponential likelihood ratio for peaks over threshold has proven a useful and efficient tool in diagnosing tail behavior and distinguishing between exponentially- and heavy-tailed families of PDFs.

Results reported here confirm that most of the precipitation exceedance probabilities follow the Pareto law and therefore, most precipitation PDF tails follow a power law (i.e. are heavy), not exponential. Threshold exceedances at stations that may have approximately exponential tails can also be thought of as limiting, “special cases" of the Pareto model, i.e. Pareto with a very large tail parameter, $\boldsymbol{\alpha}$. As $\boldsymbol{\alpha}$ increases, the Pareto distribution gets close to the exponential and the L statistics may not be sensitive enough for high values of $\boldsymbol{\alpha}$. But such cases appear to be rare in North America and, in practice, the heavy tail is clearly superior when all seasons' data is considered. The divergence from exponential is a matter of degrees. The most closely exponential tails are found at stations where similar systems and consistent processes produce virtually all of the precipitation. How far from exponential the tails are depends on location and seasonality. Places that receive precipitation from a variety of systems throughout the year tend to exhibit severely non-exponential precipitation tails. Diversity is the mother of volatility.

An exercise in limiting precipitation diversity, by considering various seasons' daily data separately, indicates that daily precipitation tails are heavier when all seasons' data is included compared to individual seasons, further illustrating the idea that the diversity of precipitation-producing systems is essentially responsible for precipitation volatility. This is also clear from examining the geographical patterns of volatility in all, as well as in individual, seasons. Seasonal analyses also indicate that overall heavy tails occur only where at least one season's data classifies as heavy-tailed. This observation is consistent with theory, according to which, a mixture of distributions can yield a heavy tail, only if at least one of the distributions in the mix is heavy-tailed.

Severe diversity in topography also appears capable to fatten or flatten precipitation tails. Notably heavy precipitation tails observed along the precipitous Mexican part of the Gulf Coast in winter, heavy (light) tails observed to the southeast (northwest) of the Appalachian Range during hurricane season, and heavy (light) tails observed east (west) of the Great Continental Divide, all testify to the additional effects of topography on precipitation volatility. Orographic gradients in volatility can be explained by orographically imposed limits on the diversity of precipitation producing systems. A mountain range can enhance or reduce the precipitation produced by specific types of storms while totally or partially limiting the spatial range of some precipitating disturbances to one side of the mountain.

This work is meant as a first step to relate hydrologic weather extremes to climate. It is based on a physically and mathematically consistent attempt to choose a correct model for describing the probability of extreme precipitation events directly from the stochastic properties of the data. Our empirical results strongly suggest the superiority of heavy tails to the traditional exponential. Our approach, results, and their implications, particularly that the fundamental shape of the precipitation probability tail is geographically specific resulting from multiplicity of precipitation producing mechanisms, differ substantially with other work on this subject, most recently and notably by Wilson and Toumi (2005). We believe that the spatial structure and coherence of our results and their solid grounding in climate and weather patterns over the North American continent demonstrate the benefits of our simple approach based on letting the data show us the nature of its tails. Limiting the distributional possibilities based on the stochastic properties of the data is a positive step forward because the scientific potential and practical implications of classifying precipitation probability tails into heavy- or exponentially-tailed families are considerable.

The real prospects of climatic change emphasize the need for mathematical models of extremes consistent with reality. The urgency of this problem is amplified by the prospect that global hydrologic change may disproportionately manifest itself in increased frequency of extreme precipitation. This view is supported by theoretical reasoning (Allen and Ingram 2002; Trenberth et al. 2003; Karl and Trenberth 2003) and climate pojections (e.g., Zwiers and Kharin 1998; Kharin and Zwiers 2000, Semenov and Bengtsson 2002; Hegerl et al. 2004; Wehner 2004, Groisman et al. 2005) as well as empirical evidence for increasing trends in the frequency of extreme daily precipitation worldwide (e.g., Tsonis 1996, Easterling et al. 2000; Folland and Karl 2001; Groisman et al. 1999a, 2004, 2005). It is difficult, however, to define and account for high frequency precipitation extremes without a reasonable probability distribution model. The choice between heavy-tailed and exponentially-tailed models is of a qualitative nature. The heavy tailed distributions have much larger high percentiles *relative* to the rest of the data values than the exponentially tailed ones. That implies that in places where heavy tailed models are appropriate, and especially if observed and modeled trends continue, future large events may be much larger than those observed to date. We need to be prepared for such possibilities. The exponentially-tailed models of precipitation will not be able to predict very large (relative to the observed data) events, because their mathematical properties do not allow for such extremes. In a future study, we plan to apply our methodology to quantify precipitation volatility in temporally evolving observational and dynamical modeling contexts.

Wherever stochastic modeling of regional precipitation extremes is of interest, a classification of local precipitation into exponential or heavy-tailed families of PDFs is the logical first step towards choosing a specific PDF to represent local precipitation. The classification exercise presented here can be easily adapted to observations at sub-daily temporal scales, wherever available, to better focus on precipitation intensity. Daily accumulations do not represent precipitation intensity well, especially for short-lived local convective systems, and/or fast-moving disturbances. Dynamical atmospheric models are notoriously deficient in their simulation of precipitation intensity, frequency, and, thus, PDFs. Our methodology can be used as a first step to determine the extent and possible causes of dynamical model shortcomings as well as to develop statistical correction schemes for global and regional dynamical model outputs. This methodology can also be used in studies of other hydrologic or broader environmental extremes resulting from volatile processes, e.g. streamflow and wildfire severity. Wherever coherent spatial patterns in volatility are found, they can guide the estimation/interpolation of environmental risk for data-poor locations. The future is bright with interesting problems.

Acknowledgments

This work was supported through NSF grant ATM-0236898 "Modeling, variability and predictability of North American hydrologic extremes". Funding was also provided by the California Climate Change Center, sponsored by the California Energy Commission's Public Interest Energy Research Program and by the NOAA Office of Global Programs, under the California Applications Program. We have enjoyed and benefited from interesting discussions with Dan Cayan, Mike Dettinger, Kelly Redmond and Tereza Cavazos and we gratefully acknowledge their input. We thank D. Rominger for programming support. Results of this study contribute to the NATO Science for Peace project (SFP 981044) "Extreme precipitation events: their origins, predictability and impacts".

References

Allen, M.R. and Ingram, W.J. (2002) Constraints on future changes in climate and the hydrological cycle. Nature 419, 224-232.

Anderson, P.L., and Meerschaert, M.M. (1998) Modeling river flows with heavy tails. *Water Resources Research*, **34**, 2271-2280.

Balkema, A.A., and de Haan, L. (1974) Residual life time at great age. Annals of Probability 2, 792-804.

Bradley, B. O., and Taqqu, M. (2003) Financial risk and heavy tails. In: S.T. Rachev (Ed.), *Handbook of Heavy Tailed Distributions in Finance*. Elsevier Science, Amsterdam, pp. 35-103.

Davidson A.C. and Smith, R.L. (1990) Models for exceedances over high thresholds. J. Roy. Statist. Soc. Ser. B 52, 393-442.

Easterling, D.R., Evans, J.L., Groisman, P.Ya., Karl, T.R., Kunkel, K.E. and Ambenje, P. (2000) Observed variability and trends in extreme climate events: A brief review. Bull. Amer. Meteorol. Soc. 81, 417-425.

Falk, M. (1995) On testing the extreme value index via the pot-method. Annals of Statistics 23, 2013-2035.

Folland, C.K. and Karl, T.R. (2001) Observed climate variability and change. In: J.T. Houghton, et al., (Eds.), *Climate Change 2001: The Scientific Basis. Contribution of Working Group 1 to the Third IPCC Scientific Assessment.* Cambridge University Press, Cambridge and New York, pp. 99-181.

Groisman, P.Ya., Karl, T.R., Easterling, D.R., Knight, R.W., Jamason, P.F., Hennessy, K.J., Suppiah, R., Page, C.M., Wibig, J., Fortuniak, K., Razuvaev, V.N., Douglas, A., Førland, E. and Zhai, P-M. (1999a) Changes in the probability of heavy precipitation: Important indicators of climatic change. Climatic Change 42, 243-283.

Groisman P.Ya., Knight, R.W., Karl, T.R., Easterling, D.R., Sun, B. and Lawrimore, J.M. (2004) Contemporary changes of the hydrological cycle over the contiguous United States: Trends derived from in-situ observations. J. Hydrometeor. 5, 64-85.

Groisman, P. Ya., Knight, R.W., Easterling, D.R., Karl, T.R. and Hegerl, G.C. (2005) Trends in intense precipitation in the climate record. J. Climate 18, 1326-1350.

Gross, J., Heckert, A., Lechner, J. and Simiu, E. (1994) Novel extreme value estimation procedures: Application to extreme wind data. In: J. Galambos et al. (Eds.), *Extreme Value Theory and Applications*, Kluwer Academic Publishers, Dordrecht, pp.225-284.

Hegerl, G.C., Zwiers, F.W., Stott, P.A. and Kharin, V.V. (2004) Detectability of anthropogenic changes in temperature and precipitation extremes. J. Climate 17, 3683-3700.

Johnson, N. L., Kotz, S. and Balakrishnan, N. (1994) *Continuous Univariate Distributions*, 2nd edition, Wiley, New York.

Karl, T.R. and Trenberth, K.E. (2003) Modern global climate change. Science 302, 1719-1723.

Katz, R.W., Parlange, M.B. and Naveau, P. (2002) Statistics of extremes in hydrology. Advances in Water Resources 25, 1287-1304.

Kharin, V.V., and Zwiers, F.W. (2000) Changes in the extremes in an ensemble of transient climate simulations with a coupled atmosphere-ocean GCM. J. Climate 13, 3760-3788.

Kozubowski, T.J., Panorska, A.P., Qeadan, F. and Gershunov, A. (2007) Testing exponentiality versus Pareto distribution via likelihood ratio. To be submitted. Preprint available from authors.

Lehmann, E.L. (1997) *Testing Statistical Hypotheses*, 2nd edition, Springer, New York.

Miranda S. (2003) Actualizacion de la base de datos ERIC II. Final report of the project TH-0226, IMTA internal reports.

NCDC (2003) Data documentation for data set 3200 (DSI-3200): Surface land daily cooperative summary of the day. National Climatic Data Center, Asheville, NC. [www.ncdc.noaa.gov/pub/data/documentlibrary/tddoc/td3200.pdf].

Pickands, J. (1975) Statistical inference using extreme value order statistics. Annals of Statistics 3, 119-131.

Samorodnitsky, G. and Taqqu, M. (1994) *Stable non-Gaussian Random Processes.* Chapman & Hall, New York.

Semenov V.A. and Bengtsson, L. (2002) Secular trends in daily precipitation characteristics: Greenhouse gas simulation with a coupled AOGCM. Climate Dynamics 19, 123-140.

Smith, J.A. (1989) Regional flood frequency analysis using extreme order statistics of the annual peak record. Water Resources Research 25, 311-317.

Smith, R.L. (1987) Estimating tails of probability distributions. Annals of Statistics 15, 1174-1207.

Smith, R.L. (1994) Multivariate threshold methods. In: J. Galambos et al. (Eds.), *Extreme Value Theory and Applications*, Kluwer Academic Publishers, Dordrecht, pp. 225-284.

Smith, R.L. (2001) Extreme value statistics in meteorology and the environment. *Environmental Statistics*, http://www.stat.unc.edu/postscript/rs/envnotes.pdf, Chpt. 8, 300-357.

Trenberth, K.E., Dai, A., Rasmussen, R.M. and Parsons, D.B. (2003) The changing character of precipitation. Bull. Amer. Meteorol. Soc. 84, 1205-1217.

Tsonis A.A. (1996) Widespread increases in low-frequency variability of precipitation over the past century. Nature 382, 700-702.

van Montfort, M.A.J. and Witter, J.V. (1985) Testing exponentiality against generalized Pareto distribution. Journal of Hydrology 78, 305-315.

Vincent, L.A., and Gullett, D.W. (1999) Canadian historical and homogeneous temperature datasets for climate change analyses. Int. J. of Climatology 19, 1375-1388.

Wehner, M.F. (2004) Predicted twenty-first-century changes in seasonal extreme precipitation events in the Parallel Climate Model. J. Climate 17, 4281-4290.

Wilson P.S. and Toumi, R. (2005) A fundamental probability distribution for heavy rainfall. Geophys. Res. Lett. 32, L14812, doi:10.1029/2005GL022465.

Zolina, O., Kapala, A., Simmer, C. and Gulev, S.K. (2004) Analysis of extreme precipitation over Europe from different reanalyses: A comparative assessment. Global and Planetary Change 44, 129-161.

Zwiers, F.W. and Kharin, V.V. (1998) Changes in the extremes of the climate simulated by CCC GCM2 under CO2 doubling. J. Climate 11, 2200-2222.