WATER RESOURCES RESEARCH, VOL. 44, W12439, doi:10.1029/2008WR007316, 2008

Improving the simulation of extreme precipitation events by stochastic weather generators

Eva M. Furrer[1] and Richard W. Katz[2]

Received 29 July 2008; revised 22 September 2008; accepted 3 October 2008; published 27 December 2008.

[1] Stochastic weather generators are commonly used to generate scenarios of climate variability or change on a daily timescale. So the realistic modeling of extreme events is essential. Presently, parametric weather generators do not produce a heavy enough upper tail for the distribution of daily precipitation amount, whereas those based on resampling have inherent limitations in representing extremes. Regarding this issue, we first describe advanced statistical tools from ultimate and penultimate extreme value theory to analyze and model extremal behavior of precipitation intensity (i.e., nonzero amount), which, although interesting in their own right, are mainly used to motivate approaches to improve the treatment of extremes within a weather generator framework. To this end we propose and discuss several possible approaches, none of which resolves the problem at hand completely, but at least one of them (i.e., a hybrid technique with a gamma distribution for low to moderate intensities and a generalized Pareto distribution for high intensities) can lead to a substantial improvement. An alternative approach, based on fitting the stretched exponential (or Weibull) distribution to either all or only high intensities, is found difficult to implement in practice.

Citation: Furrer, E. M., and R. W. Katz (2008), Improving the simulation of extreme precipitation events by stochastic weather generators, *Water Resour. Res.*, *44*, W12439, doi:10.1029/2008WR007316.

1. Introduction

[2] Stochastic weather generators are commonly relied on to simulate daily time series of weather, including elements such as minimum and maximum temperature and precipitation amount [*Richardson*, 1981; *Wilks and Wilby*, 1999]. Such simulations are sometimes intended to reflect solely natural variations under the present climate or, alternatively, to be consistent with large-scale global change [*Katz*, 1996; *Wilks*, 1992]. These simulations are typically used as inputs in assessments of the societal impacts of variations in weather and climate, for example in statistical downscaling of the output from a general circulation model, see the study by *Wilby et al.* [1998] for more details and approaches. Therefore the realistic modeling of the frequency and severity of extreme weather events (e.g., hot or cold spells or high precipitation amounts) is essential.

[3] There are several types of stochastic weather generator, including parametric [e.g., *Richardson*, 1981] and resampling [e.g., *Rajagopalan and Lall*, 1999] approaches. The parametric approach involves a formal stochastic model for daily time series of weather variables, with parametric probability distributions, such as the normal for temperature and the gamma for precipitation "intensity" (i.e., nonzero amount). The resampling approach imposes fewer constraints about the structure of the time series of weather variables, especially no parametric assumptions about their distributions. Although both these types of weather generator perform reasonably well in terms of reproducing average characteristics of variables, neither necessarily performs particularly well in terms of simulating extremes, especially high precipitation amounts [*Sharif and Burn*, 2006; *Wilks*, 1999].

[4] Extensive efforts have been devoted, particularly within the hydrology community, to statistically modeling high precipitation amount, with much evidence of its distribution being heavy-tailed [e.g., *Koutsoyiannis*, 2004]. Ideally, stochastic weather generators ought to model precipitation extremes in a manner consistent with this information. Yet achieving a unified treatment can be difficult in practice. One advantage of stochastic weather generators is their relatively simple structure, making feasible modeling multiple variables, incorporating annual cycles, and introducing covariates, such as the El Niño-Southern Oscillation (ENSO) phenomenon [*Furrer and Katz*, 2007]. So it would be desirable to retain this simplicity in improving the performance of weather generators in simulating extremes.

[5] In the present paper, we focus on the parametric type of stochastic weather generator, in attempting to improve the simulation of high precipitation amounts. Because of its convenience for implementing such improvements, the generalized linear modeling (GLM) framework for stochastic weather generators, introduced by *Furrer and Katz* [2007], is used (see www.image.ucar.edu/~eva/GLMwgen/). Previous attempts at improvements have tended to be ad hoc, for the most part lacking any theoretical justification. Instead, we rely on the statistical theory of extreme values [*Coles*, 2001] to guide our choice of improvements. Various subtleties in this theory, especially so-called "penultimate" approxima-

[1]Institute for Mathematics Applied to Geosciences, National Center for Atmospheric Research, Boulder, Colorado, USA.

[2]Institute for Study of Society and Environment, National Center for Atmospheric Research, Boulder, Colorado, USA.

0043-1397/08/2008WR007316$09.00

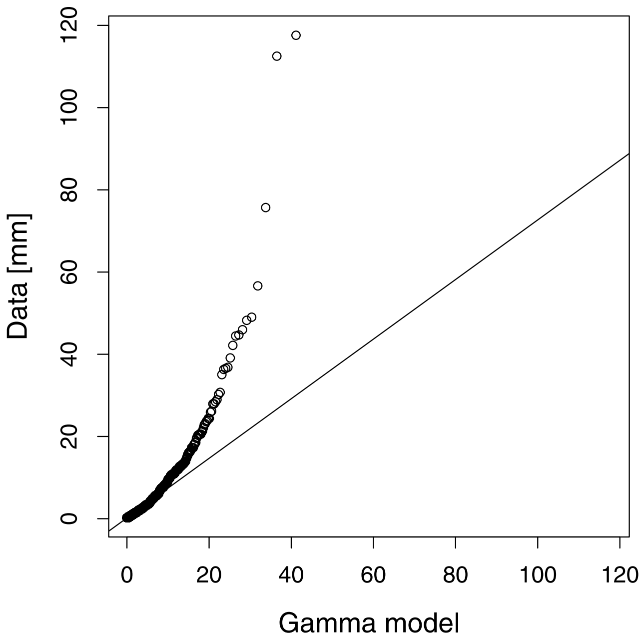

Figure 1. *Q-Q* plot of observed versus modeled gamma quantiles of precipitation intensity in July at Fort Collins.

tions, are exploited (see chapter 3 of *Embrechts et al.* [1997] and chapter 6 of *Reiss and Thomas* [2007]).

[6] Two sets of time series of daily weather are modeled, one for Fort Collins, CO for which precipitation extremes were previously analyzed by *Katz et al.* [2002] and another for Pergamino, Argentina for which minimum and maximum temperature and precipitation amount were previously analyzed by *Furrer and Katz* [2007]. First, some background on stochastic weather generators is provided, covering both parametric and resampling approaches, including previous attempts to improve the simulation of precipitation extremes (section 2). Then the modern point process approach in extreme value theory is applied to daily precipitation extremes, allowing for annual cycles and other covariates such as ENSO, as well as another distribution for high intensity, the "stretched exponential" (or Weibull), justified on the basis of a penultimate approximation (section 3). Next several approaches to improve the simulation of high precipitation amounts are considered, including a "hybrid" distribution for intensity consisting of the gamma for low to moderate values but a generalized Pareto (GP) for high values (section 4). Finally, the discussion emphasizes the remaining limitations in the various proposed approaches, particularly the difficulty in achieving parsimony (section 5).

2. Stochastic Weather Generators

2.1. Parametric Models

[7] We briefly describe the basic form of parametric stochastic weather generators, sometimes termed the Richardson model [*Richardson*, 1981]. To make wet and dry spells persist, the daily time series of precipitation occurrence is modeled with a first- or higher-order, two-state Markov chain. Given the occurrence sequence, precipitation intensities are assumed conditionally independent and identically distributed (iid). The other weather elements (e.g., minimum and maximum temperature) are linked to precipitation occurrence by shifting their conditional means and, possibly, conditional standard deviations depending on whether or not precipitation occurs. Dependence between these remaining weather elements, after adjustment for the shifts in mean and standard deviation with precipitation occurrence, is modeled as a first-order multivariate autoregressive process.

[8] Here our attention is focused primarily on how the distribution of daily precipitation intensity is modeled. If we let X_t denote the precipitation amount on the tth day, then the cumulative distribution function of intensity can be expressed as

$$F(x) = \Pr\{X_t \leq x | X_t > 0\}, \ x > 0. \tag{1}$$

In parametric generators, it is common to assume that this distribution is the gamma; that is, with probability density function given by

$$f(x;\alpha,\sigma) = \frac{1}{\sigma\Gamma(\alpha)}\left(\frac{x}{\sigma}\right)^{\alpha-1}\exp\left(-\frac{x}{\sigma}\right), \ x > 0, \alpha > 0, \sigma > 0. \tag{2}$$

Here α denotes the shape parameter and σ the scale parameter. For example, *Furrer and Katz* [2007] used the gamma distribution, with a seasonal cycle in the logarithm of its scale parameter [i.e., $\ln(\sigma)$], as well as possible dependence on ENSO, in an application of their GLM approach to daily weather at Pergamino.

[9] Figure 1 shows a quantile-quantile (*Q-Q*) plot for the fit of the gamma distribution, with parameters estimated by maximum likelihood, to daily precipitation intensity at Fort Collins only during the single month of July, 1900–1999 (see the study by *Katz et al.* [2002] for more information about this data set). This month falls within the time of the year during which floods are most likely in this region. Despite indicating an acceptable fit for low to moderate values of precipitation intensity (i.e., the overwhelming majority of the observations), the fit for higher values is rather poor. The upper tail of the gamma distribution is not heavy enough, underestimating the likelihood of high precipitation intensity to a considerable extent. As will be seen later, the distribution of intensity is not nearly as heavy-tailed at Pergamino as at Fort Collins.

[10] Because the other variables in the Richardson type of weather generator only depend on precipitation through its occurrence, it would be straightforward to implement any changes in how intensity is modeled. Proposed improvements to fitting the distribution of precipitation intensity in weather generators have included substituting another parametric distribution for the gamma, such as a mixture of two exponential distributions [*Johnson et al.*, 1996; *Wilks*, 1999]. Originally proposed by *Smith and Schreiber* [1974] as a statistical model for intensity, a mixture of two exponential distributions has a probability density function of the form

$$f(x;\omega,\sigma_1,\sigma_2) = (1-\omega)\frac{1}{\sigma_1}\exp\left(-\frac{x}{\sigma_1}\right) + \omega\frac{1}{\sigma_2}\exp\left(-\frac{x}{\sigma_2}\right),$$
$$x > 0, 0 < \omega < 1, \sigma_1 > 0, \sigma_2 > 0, \tag{3}$$

where ω denotes the mixing weight and σ_1 and σ_2 the two scale parameters. This type of mixture distribution does tend to result in a heavier tail than the gamma when fit to intensity, but still not necessarily heavy enough [*Wilks*, 1999].

[11] Alternatively, the use of a parametric form can be relaxed and the distribution of intensity fit by nonparametric techniques, such as kernel density smoothing. For example, *Semenov* [2008] used a simple binning technique. While this technique does result in acceptable performance in simulating extremes consistent with observations [*Semenov*, 2008], it has the serious limitation of not being capable of simulating extremes more than negligibly higher than those observed. Further, it would no longer be straightforward to incorporate covariates into the weather generator (e.g., seasonal cycles in intensity or dependence on ENSO).

[12] In the stochastic modeling of precipitation, but outside the realm of weather generators, several proposals for improvements in the treatment of extremes have appeared in recent years. In a conceptual stochastic model of rainfall, *Cameron et al.* [2001] used the generalized Pareto (GP), instead of the exponential, for the distribution of high intensities from a rain cell. The cumulative distribution function of the GP is given by

$$F(x;\xi,\sigma,u) = 1 - \left[1 + \xi\frac{x-u}{\sigma}\right]^{-1/\xi}, \quad x > u, 1 + \xi\frac{x-u}{\sigma} > 0, \sigma > 0. \tag{4}$$

Here ξ denotes the shape parameter, where positive ξ implies a heavy tail, negative ξ a bounded tail and the limiting case of $\xi \to 0$ a light tail (i.e., the exponential distribution), σ the scale parameter, and u the threshold for high intensity. Such an approach has a firm basis in extreme value theory [e.g., *Coles*, 2001], which implies that the upper tail of essentially any distribution must be approximately GP for sufficiently high u.

[13] Making use of a heuristic physical argument, *Wilson and Toumi* [2005] argued that the distribution of high precipitation intensity ought to be approximately of the stretched exponential (or Weibull) form. Its cumulative distribution function is given by

$$F(x;c,\sigma,u) = 1 - \exp\left[-\left(\frac{x-u}{\sigma}\right)^c\right], x > u, c > 0, \sigma > 0, \tag{5}$$

where c denotes the shape parameter, σ the scale parameter and u a threshold. This distribution can possess an apparent heavy tail, as formally established through a penultimate approximation in extreme value theory (see section 3). It is appealing for stochastic weather generators because it can be used within a GLM framework (chapter 13 of *McCullagh and Nelder* [1989]), which, for example, is not possible for the GP distribution.

[14] To avoid choosing a threshold for high intensity, *Vrac and Naveau* [2007] used a dynamic mixture with gamma and GP components. Unlike an ordinary mixture, the dynamic mixture is designed to weight the gamma more for low-intensity values, the GP more for high values. One drawback to this approach is that it requires the estimation of an unusually high number of parameters, at least six in practice.

2.2. Resampling Approach

[15] The resampling approach to stochastic weather generation involves drawing from the original observations with replacement, termed the "bootstrap" in statistics [*Efron and Tibshirani*, 1993]. In practice, the resampling scheme may be rather complicated, with the resampled objects being vectors of weather observations (e.g., daily minimum and maximum temperature and precipitation amount) and the resampling being restricted to nearest neighbors to preserve temporal dependence [*Rajagopalan and Lall*, 1999]. Several limitations, especially with respect to extremes, are inherent to any resampling scheme. Since only observed values can be resampled, in particular, no value in between the second highest value and the highest value is possible, etc. Moreover, a precipitation intensity higher than the maximum observed value can never be generated. These limitations of the resampling approach have been previously recognized [e.g., *Rajagopalan and Lall*, 1999]. Some sort of smoothed bootstrap would rectify the issue of not being possible to generate values between those observed. For example, *Sharif and Burn* [2006] introduced smoothness through perturbing the output of the resampling algorithm by adding a noise term. However, it is not clear how to modify the resampling approach to generate values higher than the observed maximum, other than in an ad hoc fashion.

3. Upper Tail Modeling

[16] In principle we are interested in the simulation from the entire distribution of precipitation intensity, which traditionally is deficient for high intensities. Our goal being to improve on the simulation of these high intensities, it is crucial to first understand the upper tail behavior alone in order be able to make plausible proposals for improvement. The statistical modeling of high precipitation intensity is naturally linked to extreme value theory, which focuses on the behavior and the modeling of the upper tails of distributions. In section 4 we will show how to use the obtained results for extremes in a unified modeling of the distribution of precipitation intensity.

3.1. Extreme Value Analysis

[17] We begin this section by explaining briefly how to characterize extreme value behavior using a two-dimensional point process approach, see the work of *Smith* [1989], chapter 5 in the book by *Leadbetter et al.* [1983], and chapter 7 in the book by *Coles* [2001] for more details. One of the advantages of this representation is its unification of the more traditional block maxima (generalized extreme value distribution, GEV) and peaks-over-threshold (POT, GP distribution) approaches. It uses the GEV parameterization given by the cumulative distribution function

$$F(x;\xi,\sigma^*,\mu) = \exp\left\{-\left[1 + \xi\frac{x-\mu}{\sigma^*}\right]^{-1/\xi}\right\}, \; 1 + \xi\frac{x-\mu}{\sigma^*} > 0, \tag{6}$$

with location parameter $-\infty < \mu < \infty$, scale parameter $\sigma^* > 0$ and shape parameter ξ having the same interpretation as for the GP (with the limiting case of $\xi \to 0$ being the Gumbel distribution). The scale parameters of the GEV and

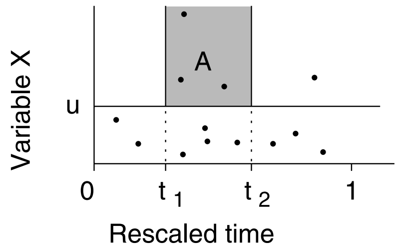

Figure 2. Illustration of point process representation of extremes with event $A = [t_1, t_2] \times (u, \infty)$.

GP parameterizations, σ^* and σ respectively, are related through $\sigma = \sigma^* + \xi(u - \mu)$. The point process approach is formulated in terms of the GEV parameters, which have the advantage of not depending on a threshold, while not reducing the data to block maxima and while including the threshold exceedance rate in the inference. As a consequence, non-stationarity can be easily and naturally introduced through covariate effects in the parameters and, theoretically, there is no difficulty in working with time-varying thresholds.

[18] For the modeling of extreme values, a Poisson process is a natural model in a limiting sense, as seen in the following. We start with the simplified situation of a sample of iid random variables $X_1,\ldots, X_n$ with common cumulative distribution function F. Note that, in practice, the iid assumption is often not met implying the need to decluster temporally close excesses over the threshold u [*Coles*, 2001]. Then the distribution of the appropriately normalized maximum $M_n = \max\{X_1,\ldots, X_n\}$ converges to the GEV distribution, and the distribution of the excesses over a high threshold u is approximated by a GP distribution under mild conditions on F. The sequences used to appropriately normalize the maximum, which ensure a non-degenerate limiting distribution, are not known in practice. Fortunately, if the appropriately normalized maximum has asymptotic GEV distribution, so does the maximum itself but with different location and scale parameters. Hence, for statistical applications of extreme value theory, it is customary to absorb the normalizing sequences into the estimated parameters. We use this convention in our description of the point process approach. Defining a sequence of two-dimensional point processes N_n by

$$N_n = \{(i/(n+1), X_i) \ : \ i = 1, \ldots, n\}, \tag{7}$$

it can be seen that the number of points $N_n(A)$ contained within the set $A = [t_1, t_2] \times (u, \infty)$, with $[t_1, t_2] \subset (0, 1)$ (see Figure 2) for sufficiently large u, is binomially distributed with number of trials n and success probability approximately given by

$$p \approx \frac{1}{n}(t_2 - t_1)\left[1 + \xi\left(\frac{u-\mu}{\sigma^*}\right)\right]^{-1/\xi}, \tag{8}$$

where we used the GEV/GP approximation mentioned above. Standard convergence of a binomial distribution to a Poisson limit intuitively leads to a limiting Poisson process N with intensity measure $np = (t_2 - t_1)[1 + \xi(u - \mu)/\sigma^*]^{-1/\xi}$. We refer to mathematical details in the book by *Coles* [2001], for example Theorem 7.1.1, and the references therein. Estimation of the parameters in the point process representation is by maximum likelihood and requires specialized numerical techniques in the non-stationary case. *Coles* [2001] also discusses how the simpler Poisson-GP approach, commonly used in hydrology [e.g., *Madsen et al.*, 1997], is related to the point process representation.

[19] *Katz et al.* [2002] applied the point process approach to extreme value analysis of the Fort Collins daily precipitation data over the entire year. Using a constant threshold of u = 10 mm, they found a marked annual cycle in the extreme precipitation with a peak in July, modeled with a sine wave for the location parameter μ and another sine wave for the transformed scale parameter $\ln\sigma^*$. Fairly strong evidence of a heavy tail was found, with an estimate of ξ of about 0.18 under the constraint of it having no annual cycle. Note that little evidence pointed to dependence within the excesses over u and therefore no declustering was applied in this case.

[20] Our goal is to apply the point process approach to the daily precipitation data from Pergamino over the entire year (same data set analyzed by *Furrer and Katz* [2007], 63 years of data during 1932–2003). As for the POT approach, the first thing is to choose a threshold and, since for a weather generator a model is needed for the entire year, we need a year-round threshold as well. A time-varying threshold is theoretically possible and can be used to account for seasonal variation ensuring enough observed excesses in each season. However, because time-varying thresholds can lead to numerical instabilities in the maximization of the likelihood, it is safer to start with a fixed threshold, avoiding potential numerical difficulties and providing a good reference point if a time-varying threshold should be necessary.

[21] Figure 3 shows precipitation intensity data at Pergamino. Only the last ten years of the record period are shown for visibility reasons. Intensities higher than a constant threshold of u = 16 mm are highlighted, as well as distinguishing between the winter and summer half of the year. These data (and the data from the rest of the record period, not shown here) do not indicate that a time-varying threshold is necessary and, therefore, we will account for seasonality only by considering continuously varying parameters. A common tool for threshold selection is to fit the model repeatedly using a set of candidate thresholds, and concentrating on a range of thresholds for which the resulting parameter estimates do not change too much. Since we need a year-round threshold, it should not be too high in order to have sufficient observations during the drier winter months. The selected threshold is a good compromise with respect to these points. Again, little evidence pointed to the need for declustering.

[22] Table 1 contains parameter estimates and Bayesian information criterion (BIC) values [*Schwarz*, 1978] for several models fitted to the Pergamino precipitation data using a fixed threshold of u = 16 mm. The models include a seasonal cycle and an ENSO effect (same ENSO index as used by *Furrer and Katz* [2007]) in the location parameter μ and in the logarithm of the scale parameter $\ln\sigma^*$ (i.e., same form of annual cycle as in the study by *Katz et al.* [2002] applied to the Fort Collins data). No covariates are introduced for the shape parameter ξ, in order to keep the model simple and easily interpretable. Separate fits of a point process model to winter and summer precipitation data

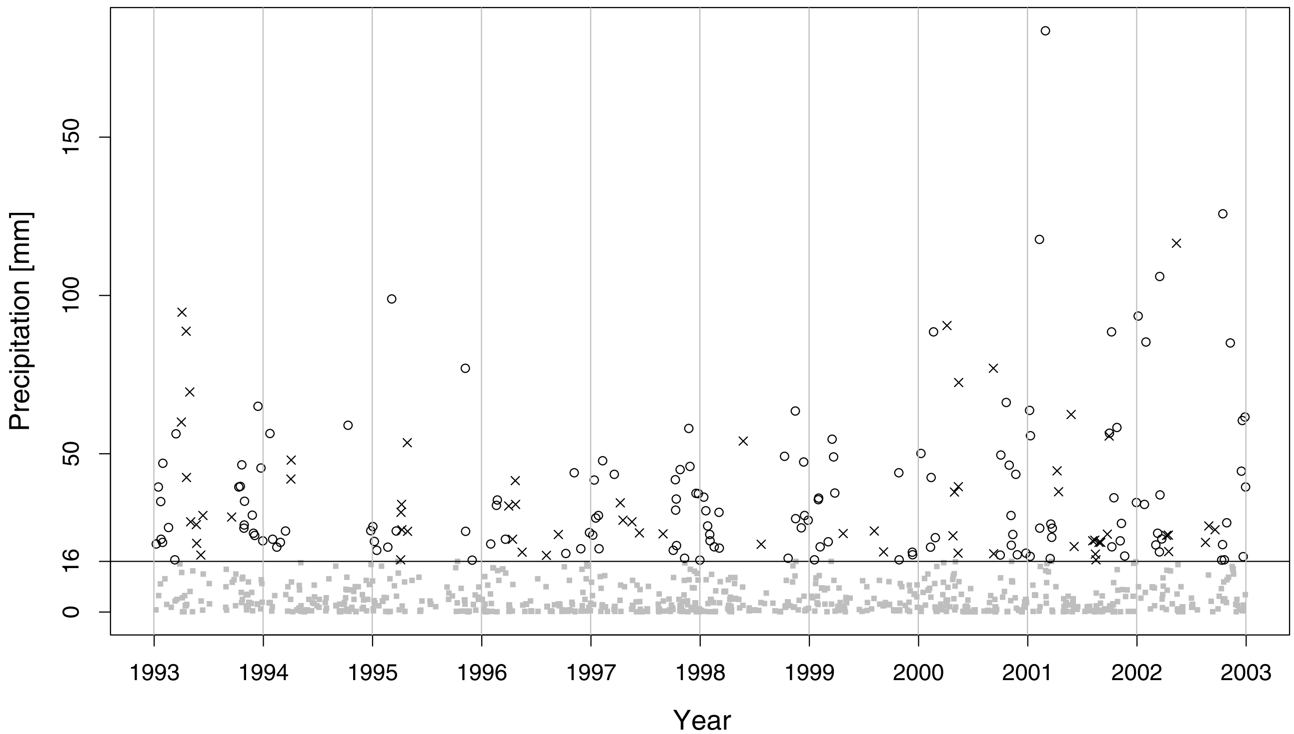

Figure 3. Precipitation intensity from the years 1993 to 2002 at Pergamino. Values more than 16 mm during the southern hemisphere winter months (April-September) are marked by x's; those of the summer months (October-March) are marked by circles.

suggest that the shape parameter might be smaller during the drier winter months, but we did not explore this any further. Model selection within the listed models, on the other hand, does not change much if considering winter and summer separately, justifying our year-round application. The model with minimum BIC includes a seasonal cycle and an ENSO effect in the location parameter and a seasonal cycle in the logarithm of the scale parameter, while the shape parameter is constant. The estimated parameters indicate a higher location parameter, i.e., higher precipitation extremes, as well as a higher scale parameter, i.e., higher variability in precipitation extremes, in southern hemisphere summer and also a higher location parameter for positive values of the ENSO index (i.e., during El Niño events). The magnitude of the shape parameter estimates are rather small (compared to those for Fort Collins), ranging from about 0.05 for models without seasonal cycle in the logarithm of scale to about 0.10 for several alternative models.

[23] For the construction of diagnostic plots, *Coles* [2001] suggests converting the point process parameter values to the corresponding GP parameters and using those to transform the observed excesses to standard exponential variables, producing a *Q-Q* plot on an "exponential" scale. Figure 4 (left) shows such a *Q-Q* plot for the chosen model. The point process model is able to capture the heavy tail of precipitation excesses in Pergamino resolving the problems of the gamma model fit to all intensities (i.e., not just excesses over a high threshold) [*Furrer and Katz*, 2007] as seen in Figure 4 (right).

3.2. Penultimate Approximations

[24] The GEV/GP distributions in extreme value theory can be viewed as "ultimate" approximations. A lesser known concept involves more refined penultimate approximations. Considering block maxima M_n of finite length n, for which the shape parameter of the asymptotic (ultimate) GEV distribution equals 0, i.e., the Gumbel distribution, a more accurate approximation is actually obtained by using a

Table 1. Estimated Parameters and BIC Values (Minimum in Bold) for Candidate Point Process Models for Daily Precipitation Extremes Over the Entire Year With a Threshold of 16 mm at Pergamino

Location μ				Scale $\ln\sigma^*$					
Constant	Cosine	Sine	ENSO	Constant	Cosine	Sine	ENSO	Shape ξ	BIC
76.52				22.31				0.072	5032.92
74.36	9.76	2.94		21.16				0.044	4869.20
78.07				9.76	2.94	21.16		0.063	4901.39
74.83	25.25	6.81		3.10	0.27	0.07		0.101	4859.30
74.20	9.76	2.81	1.94	21.16				0.044	4867.66
74.62	25.28	6.61	1.87	3.11	0.27	0.07		0.103	**4858.24**
75.10	25.92	6.73	1.85	3.12	0.28	0.07	$<10^{-2}$	0.105	4868.38

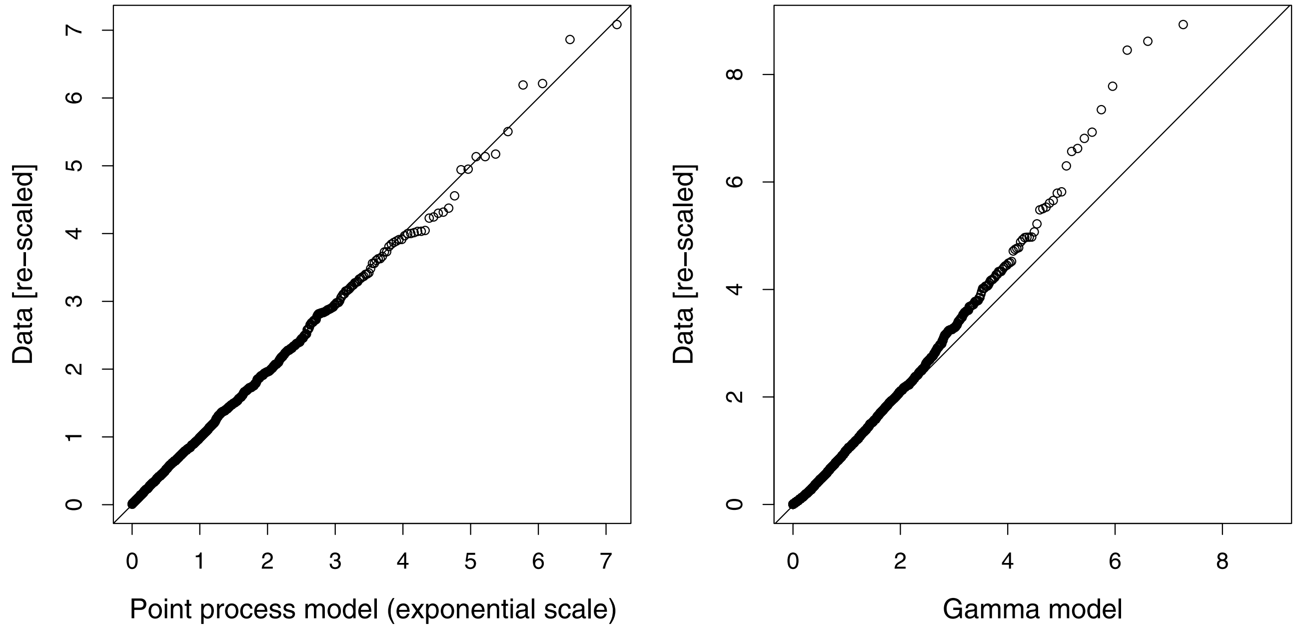

Figure 4. (left) *Q-Q* plot of observed versus modeled GP quantiles of precipitation excesses over 16 mm and (right) *Q-Q* plot of observed versus modeled gamma quantiles of precipitation intensity for the entire year at Pergamino.

non-zero shape parameter, ξ_n say. In statistical terms this means that no advantage is obtained by constraining the estimation procedure to $\xi = 0$ even if it is known that the correct limiting distribution is the Gumbel. This approximation applies equally well to the corresponding GP distribution for the excess over a high but finite threshold. Precise conditions and related theoretical results can be found in chapter 3 of *Embrechts et al.* [1997] and chapter 6 of *Reiss and Thomas* [2007]. Examples of the use of penultimate approximations in the geophysical sciences and related literature include that of *Abaurrea and Cebrián* [2002] for drought analysis and that of *Cook and Harris* [2004] for extreme winds in engineering design.

[25] For our purposes it is sufficient to provide an expression for ξ_n together with a heuristic argument for how to obtain it. In the notation of section 3.1 and assuming additionally that the first and second derivatives of F exist and that F has an infinite right endpoint, we have that the distribution of M_n can, on the one hand, be approximated by a GEV distribution with parameters ξ, σ^*, μ and, on the other hand, we have that

$$\Pr\{M_n \leq x\} = [F(x)]^n \approx \exp\{-n[1 - F(x)]\} \tag{9}$$

for n and x large such that $n[1 - F(x)] \approx$ constant. Equating these two approximations, substituting $x = \mu$, differentiating and some further algebra yields

$$\xi_n = \frac{-F''(\mu_n)}{n[F'(\mu_n)]^2} - 1 = \left(\frac{1}{H}\right)'(x)\big|_{x=\mu_n}$$
$$\text{where} \quad \mu_n = F^{-1}\left(1 - \frac{1}{n}\right), H(x) = \frac{F'(x)}{1 - F(x)} \tag{10}$$

and H is called the hazard function. Moreover, in all practical cases, $(1/H)'(x) \rightarrow \xi$ as $x \rightarrow \infty$. In other words, ξ_n as given in (10) can be viewed as reflecting preasymptotic behavior of the shape parameter.

[26] The stretched exponential distribution, proposed by *Wilson and Toumi* [2005] as a model for high precipitation intensity (see section 2), is not heavy-tailed in an ultimate sense, that is, the Gumbel distribution is the correct asymptotic model. However, in a penultimate sense, this distribution can have either an apparent heavy or bounded tail, depending on the value of its shape parameter c. In particular, the GEV shape parameter in the penultimate approximation for block maxima of length n from a stretched exponential can be expressed as

$$\xi_n = \frac{1 - c}{c \ \ln n}, \tag{11}$$

which is easily shown by substituting the stretched exponential cumulative distribution function (5) and its first and second derivatives into (10). It is obvious from formula (11) that c needs to be smaller than one in order for the tail to be heavy in a penultimate sense, i.e., $\xi_n > 0$. Note that $\xi_n \rightarrow 0$ as $n \rightarrow \infty$ (i.e., consistent with the ultimate approximation). *Wilson and Toumi* [2005] argued that c should universally be 2/3, but they did not actually fix it to that value in fitting the stretched exponential distribution to station daily precipitation data from all over the world. If $c = 2/3$ and $n = 100$ (e.g., precipitation occurs on about 27% of the days within a year), then $\xi_{100} \approx 0.11$, a reasonable value for the shape parameter of annual maxima of daily precipitation. As a rough approximation for the frequency of occurrence of precipitation within a year, we will take $n = 100$ days in subsequent use of (11). Because ξ_n only varies slowly with n, this approximation is not very restrictive.

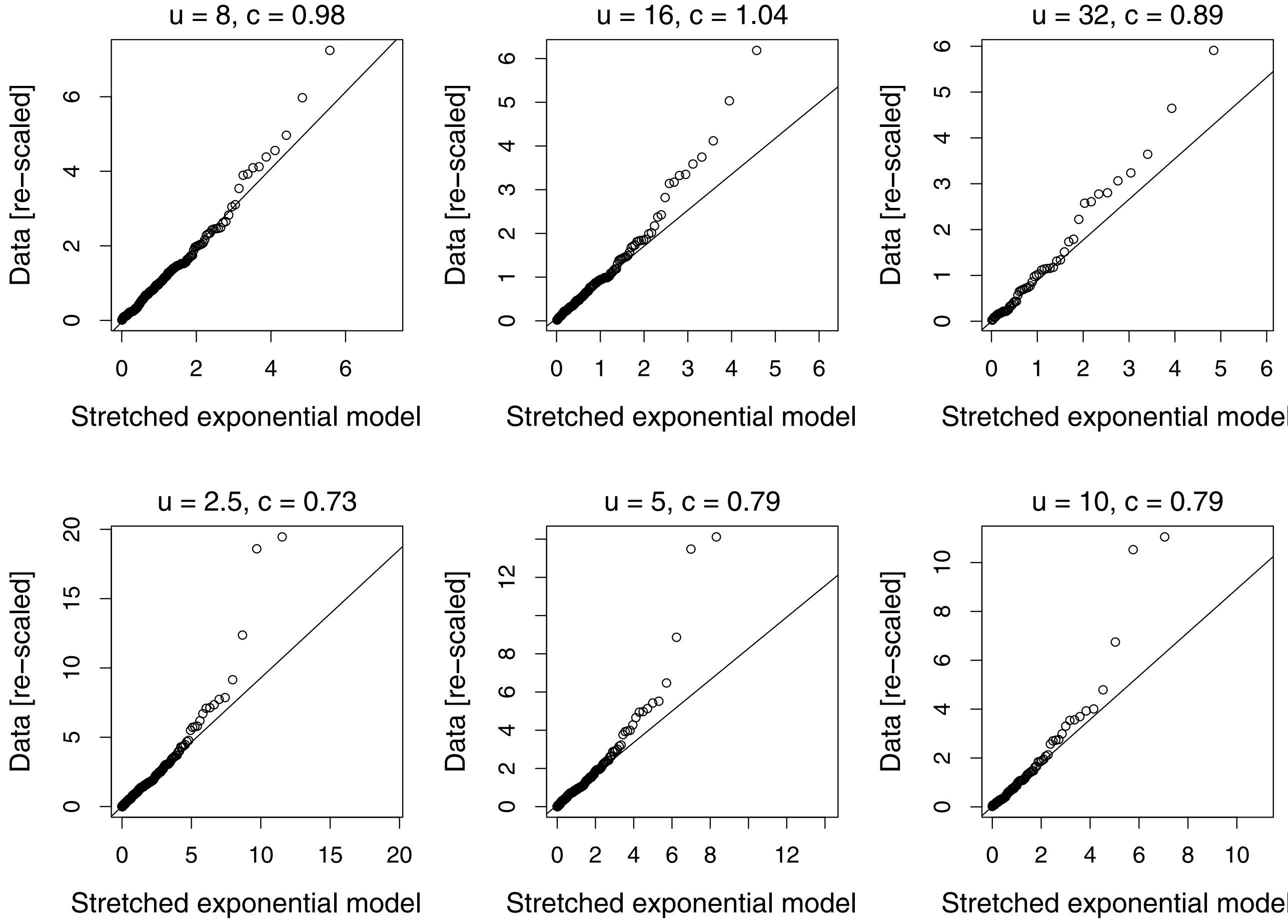

Figure 5. *Q-Q* plots of observed versus modeled stretched exponential quantiles of precipitation excesses in (top) February at Pergamino and (bottom) July at Fort Collins using different thresholds (threshold value and estimated shape parameter are indicated in the individual panels).

[27] For purposes of statistical modeling of extreme precipitation based on station data, we need to choose a threshold above which the stretched exponential is a good approximation. Because of dealing with a large number of sites, *Wilson and Toumi* [2005] used an automatic threshold selection rule based on a high quantile. We examine the choice of threshold more closely. Threshold selection is even more difficult than in the POT approach with the GP distribution, because the parameters should continue to gradually change rather than stabilizing as the threshold increases. Figure 5 shows *Q-Q* plots of stretched exponential model fits along with shape parameter estimates by maximum likelihood for excess daily precipitation intensity only from the single months of February in Pergamino and July in Fort Collins using different candidate thresholds. It indicates the difficulty of threshold choice in a stretched exponential model. First, we note that shape parameter estimates are far from the hypothesized value of 2/3 and relatively close to one (e.g., $c \approx 0.8$ is needed to produce $\xi_{100} \approx 0.05$ as is needed for Pergamino) and they do not decrease with increasing threshold. Second, the model fit evaluated in terms of the *Q-Q* plot does not improve significantly nor monotonically with increasing threshold. These results typify the difficulties that arise as well for the remaining months at the same two stations.

[28] At least for our examples, we conclude that although the stretched exponential distribution has some appealing features as a candidate model for extreme precipitation, it has the drawback of difficult threshold selection. Too low a threshold does not provide an adequate fit, too high a threshold loses the benefit of the penultimate approximation.

4. Unified Modeling

[29] Both extreme value theory approaches discussed in the previous section cannot be used per se in a weather generator framework: they model only extreme precipitation and give no indication of the behavior of the bulk of the data. In this section, we describe some possibilities of how to achieve unified modeling of the entire precipitation intensity distribution, in part driven by the results from upper tail modeling.

4.1. Single Distribution

[30] Ideally, a distribution replacing the gamma should be able to model the entire range of precipitation events, i.e., low to moderate as well as extreme events, satisfactorily while still being admissible in a GLM framework to allow the straightforward introduction of covariates and systematic assessment of uncertainties. A candidate distribution remains the stretched exponential (i.e., with no threshold),

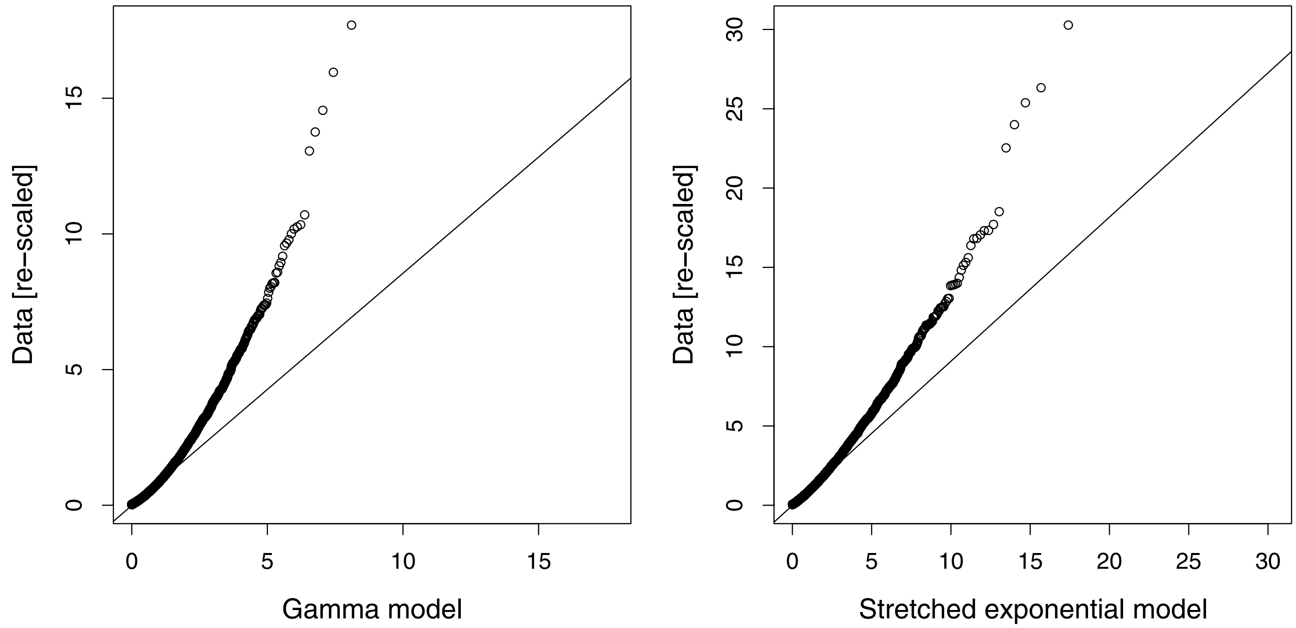

Figure 6. *Q-Q* plots of (left) observed versus gamma and (right) stretched exponential GLM quantiles of precipitation intensity for the entire year at Fort Collins.

with extremal properties discussed in the previous section, see for example chapter 13, section 3.2 of *McCullagh and Nelder* [1989] for its admissibility in a GLM framework and *Yan et al.* [2006] for a climate application. In a few instances, the stretched exponential has been previously fitted to the entire distribution of precipitation intensity, particularly for accumulations over time periods shorter than a day [*Wilks*, 1989; *Wong*, 1977]. Because the gamma distribution with shape parameter in the typical range for daily intensity (i.e., $0.5 < \alpha < 1$) cannot produce more than a non-negligible apparent heavy tail in a penultimate sense, it is conceivable that its replacement with the stretched exponential could still result in an improved fit for high intensity.

[31] Similarly to the gamma GLM, we model the logarithm of the scale of the stretched exponential distribution as a function of covariates in a GLM framework while the shape parameter is kept constant. Fitting gamma and stretched exponential GLMs to Fort Collins precipitation intensity over the entire year with seasonal cycles in the logarithms of the respective scale parameters results in the *Q-Q* plots of Figure 6. In order to produce these plots, to adjust for the non-stationary form of the fitted model, the precipitation intensity data have been re-scaled by the appropriate value of the modeled scale parameters as functions of the day of the year and then compared with respective theoretical distributions of unit scale and estimated shape parameter. As a consequence of the re-scaling, which is different for the two models but necessary to obtain *Q-Q* plots in a non-stationary setting, the two plots can only be compared qualitatively. See Table 2 for estimated coefficients and BIC values of gamma and stretched exponential models with and without seasonal cycles. The estimate of the shape parameter of the stretched exponential ($\hat{c} = 0.77$) is not too far from 2/3 but the apparent heavy tail characterized by the corresponding $\xi_{100} \approx 0.06$ is still rather light. From Figure 6 it is evident that the stretched exponential model only results in a very slight improvement (if any) over the gamma for the fit in the upper tail of the precipitation intensity distribution. At Pergamino, the estimated shape parameter of the stretched exponential and the penultimate shape parameter of the GEV take similar values ($\hat{c} = 0.70$ and $\xi_{100} \approx 0.09$) to those at Fort Collins, again a *Q-Q*-plot (not shown here) reveals a slightly better fit than that for the gamma model, but still inadequate in the upper tail.

4.2. Hybrid Approach

[32] A simple procedure to generate potentially improved higher precipitation intensity is to replace the values drawn from a gamma distribution which fall above a given threshold by values drawn from a GP distribution. However, this hybrid approach can create a problem of the underlying density function not being continuous at the threshold u. We

Table 2. Estimated Parameters and BIC Values (Minimum in Bold) for Candidate Gamma (Left) and Stretched Exponential (Right) Models for Daily Precipitation Intensity Over the Entire Year at Fort Collins

Scale $\ln\sigma$					Scale $\ln\sigma$				
Constant	Cosine	Sine	Shape α	BIC	Constant	Cosine	Sine	Shape c	BIC
1.56			0.69	41,082.34	1.36			0.76	40,365.52
1.47	−0.33	0.04	0.71	**40,845.63**	1.29	−0.30	0.05	0.77	**40,206.74**

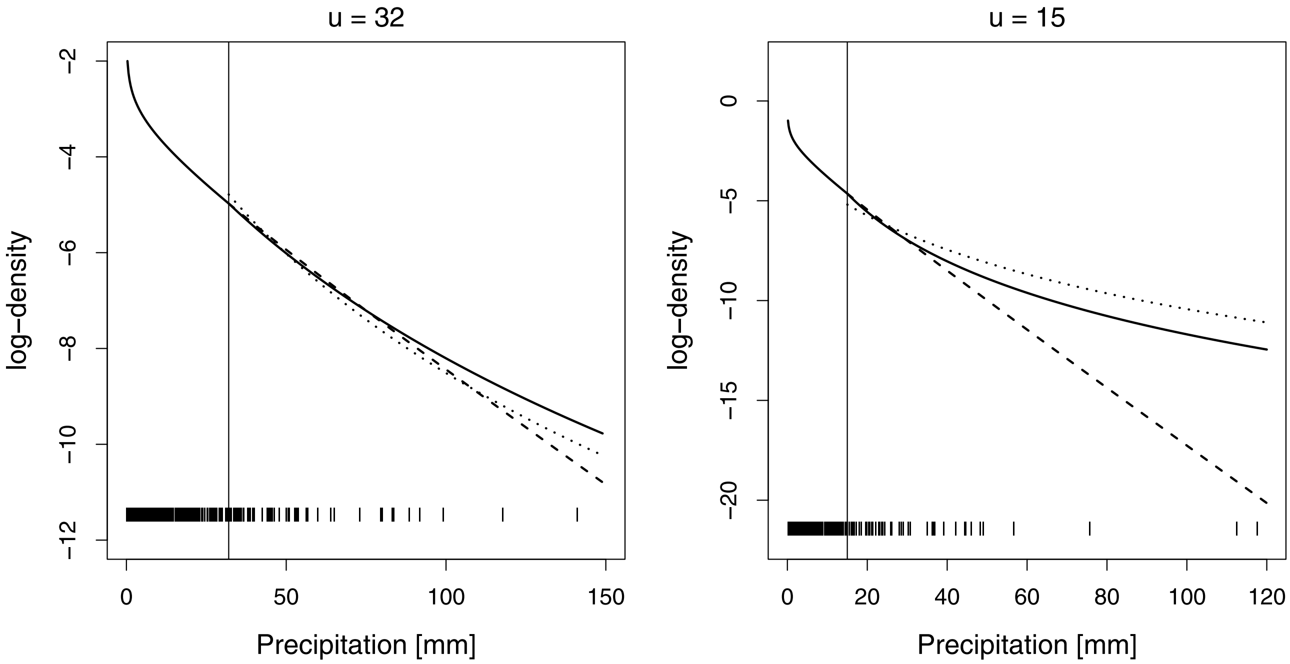

Figure 7. Log-density functions fitted to (left) February precipitation intensity at Pergamino and (right) July precipitation intensity at Fort Collins. Gamma (dashed lines), GP (dotted lines) and hybrid gamma/GP (solid lines) models are shown. The data are indicated by horizontal marks and the threshold u by a horizontal line.

refine this simple procedure through a compromise by estimating a gamma distribution from all the data (i.e., as in a conventional weather generator), estimating a GP distribution from the data above the threshold (i.e., as in a conventional extreme value analysis such as the point process technique described in section 3), and then using only the estimated shape parameter ξ of the GP while adjusting its scale parameter in order to achieve a continuous density. We call the resulting distribution a hybrid gamma/GP distribution. With this approach, we need to abandon a simple GLM setting but, as elaborated below, it is still feasible to include covariates.

[33] An expression for the adjustment of the scale parameter of the GP is obtained as follows. Denote by h the hybrid gamma/GP density we want to specify, by f the density and by F the cumulative distribution function of the fitted gamma distribution, and by g the density of the GP distribution for the excesses over the threshold u with scale parameter σ and shape parameter ξ. Observe that $g(u) = 1/\sigma$ and that the mass of g is concentrated on the interval (u, ∞) [see (4)]. By design the hybrid density satisfies

$$h(x) = f(x) \quad \text{for} \quad x \leq u \quad \text{and}$$
$$h(x) = [1 - F(u)]g(x) \quad \text{for } x > u, \tag{12}$$

where the factor $[1 - F(u)]$ ensures that h is normalized. For h to be continuous at u it is necessary that $f(u) = [1 - F(u)]g(u) = [1 - F(u)](1/\sigma)$, i.e.,

$$\sigma = \frac{[1 - F(u)]}{f(u)}. \tag{13}$$

Note that σ is hence the reciprocal of the hazard function [see (10)] of the gamma distribution evaluated at the threshold u. Further note that (13) actually holds for the hybrid of any density with the GP, not just the gamma.

[34] To apply (13) the gamma distribution needs to be estimated based on the entire data set in order to obtain an estimate of the probability of intensity higher than the threshold u. Covariate models (if any) are fitted separately for the gamma and the GP distributions as in the usual GLM and extreme value analysis settings. However, since the scale parameter of the GP distribution is adjusted using the fitted gamma distribution, its functional dependence on the covariates is no longer determined by the GP covariate model. A similar approach involving the gamma and the GP distribution is described by *Vrac and Naveau* [2007], in particular their limiting case of $\tau = 0$ (i.e., the parameter that governs the speed of the transition from gamma to GP in the weight function for the dynamic mixture) models precipitation intensity with a gamma below a threshold and a GP above a threshold. Drawbacks to the use of their approach in a weather generator setting are the discontinuity at the threshold in the limiting case and the difficulty of incorporating covariates.

[35] Figure 7 shows the fitted gamma, GP and hybrid gamma/GP log-densities for precipitation intensity from Pergamino in the single month of February (u = 32 mm) and Fort Collins in the single month of July (u = 15 mm, the logarithm of the density has been taken to improve visibility). Figure 7 shows that the hybrid density is indeed continuous and has a heavier tail than the corresponding gamma density (estimates of the GP shape parameter are 0.20 for February in Pergamino and 0.29 for July in Fort Collins). It also indicates that the heaviness of the upper tail of the Fort Collins data is much more pronounced than for

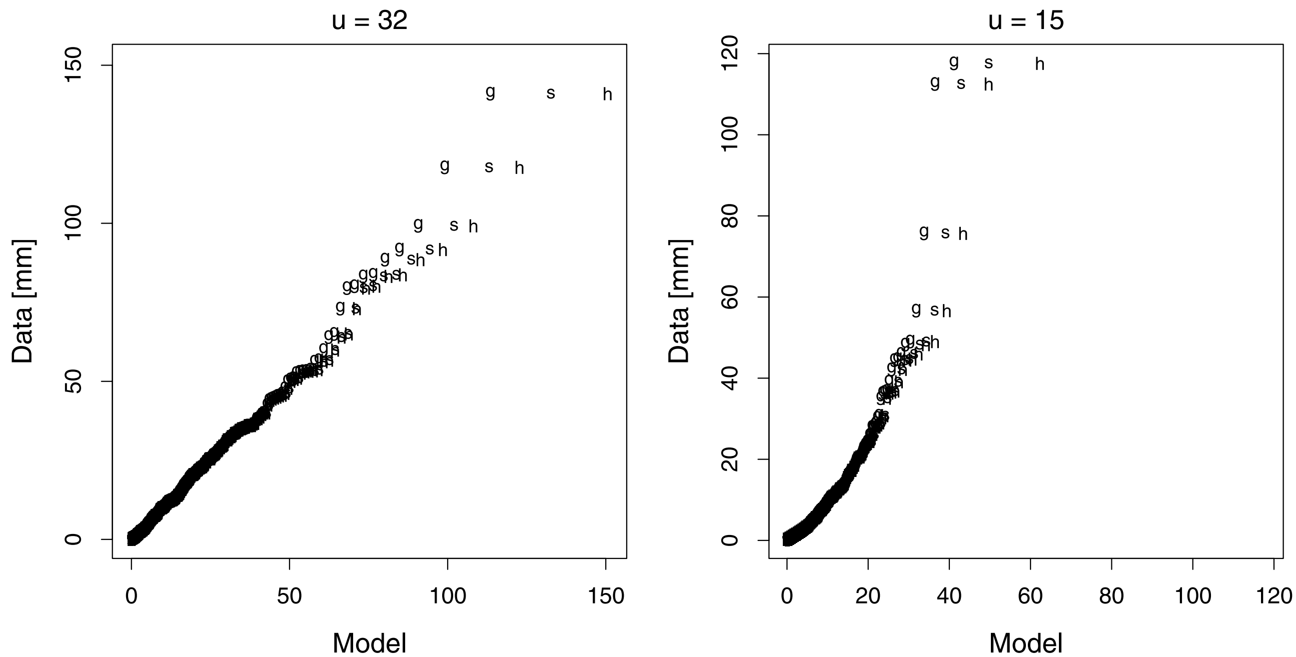

Figure 8. *Q-Q* plots of observed versus modeled gamma (*g*), hybrid gamma/GP (*h*) and stretched exponential (*s*) quantiles of precipitation intensity in (left) February at Pergamino and (right) July at Fort Collins.

Pergamino (at least in July and February, respectively), therefore more improvement is to be expected in applying this approach to the Fort Collins data. Note that the thresholds of 32 and 15mm, higher than when dealing with the entire year, were chosen because February and July are in the respective wet seasons.

[36] Figure 8 shows *Q-Q* plots of fitted gamma, hybrid gamma/GP as well as the stretched exponential models, again for precipitation intensity from Pergamino in February and Fort Collins in July. We compare the methods for a single month due to the scaling issues mentioned above which imply that comparisons of models for the entire year via *Q-Q* plots are difficult. As expected, the improvement of the hybrid approach is more clear-cut for the Fort Collins data, where it produces higher quantiles than do the stretched exponential and the gamma. None of the methods provides high enough quantiles as can be seen, for example, by the small fraction of the scale of the data quantiles that is covered by the modeled quantiles. For Pergamino both the stretched exponential and the hybrid approach produce quantiles that approximately cover the scale of the data quantiles, the largest quantile from the hybrid approach being even higher than the largest data quantile. From Figure 8 it is obvious that the hybrid approach implies a clear improvement for high intensities for cases like Fort Collins with a marked heavy tail. In less pronounced cases, as Pergamino, it does not worsen the situation.

[37] Returning to modeling precipitation over the entire year, we fit a gamma GLM with a seasonal cycle in the logarithm of the scale parameter to the precipitation intensity data from Fort Collins. A fixed threshold of u = 10mm throughout the year and a seasonal cycle in the logarithm of the scale parameter are used to estimate the shape parameter of the GP distribution ($\hat{\xi} = 0.2$), while the scale parameter is actually derived indirectly from the fitted gamma distribution (note that slightly different estimates of ξ may be produced by the point process approach than by only modeling the excesses, because the functional form of the

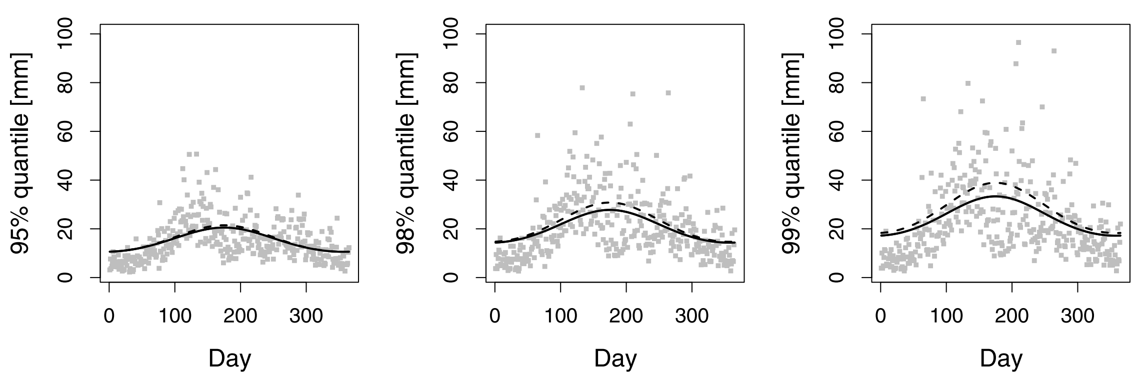

Figure 9. High quantiles of fitted gamma (solid lines) and hybrid gamma/GP (dashed lines) distributions as functions of the day of year for Fort Collins precipitation intensity. Empirical quantiles calculated from the 100 years of observed precipitation are shown as dots.

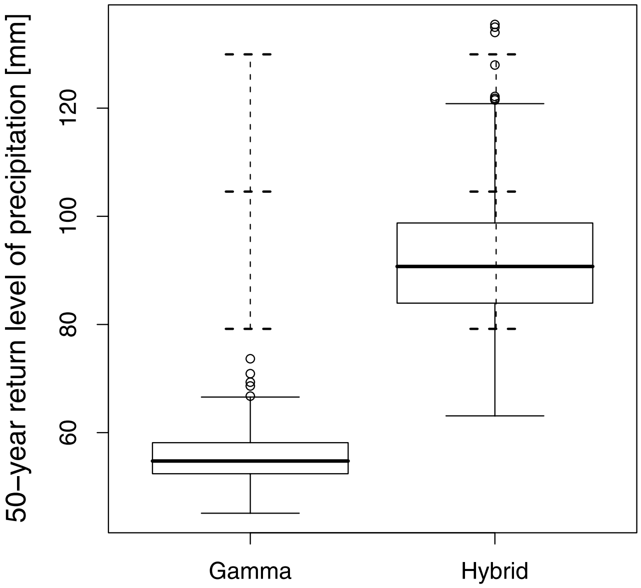

Figure 10. Boxplots of 50-year return levels of summer (April-September) maximal precipitation intensity for Fort Collins from 500 simulated samples using the gamma and the hybrid gamma/GP models along with the corresponding observed value and confidence interval (indicated by horizontal dashed lines, range of confidence interval by vertical dashed line).

dependence of the extremal parameters on the covariates is not necessarily identical). Figure 9 shows the 95%, 98% and 99% quantiles of the fitted gamma and hybrid gamma/GP distributions as functions of the day of the year, along with corresponding empirical quantiles calculated from the 100 years of observed precipitation. It indicates that the higher the quantile the more noticeable the effect of the hybrid approach, especially during the wet summer season.

[38] In order to further illustrate this ability of the hybrid approach to model high quantiles, we simulated time series of daily precipitation from a GLM weather generator using an identical (except for the lack of ENSO as a covariate) form of model for precipitation at Fort Collins as proposed by *Furrer and Katz* [2007] for Pergamino. 500 samples of 100 years of daily data have been generated, and from those summer (April-September) maxima have been calculated and GEV distributions fitted. The resulting 50-year return levels along with corresponding data values (including confidence intervals obtained by normal approximation using an approximate variance of the estimated quantile derived by the delta method [*Coles*, 2001]) are shown in Figure 10. Although still not perfect, the hybrid distribution is a clear improvement over a gamma distribution for the Fort Collins data. A similar exercise for the Pergamino data leads to a less tangible improvement. However, as remarked before, the heaviness of the tail for Pergamino precipitation is less pronounced and even if a hybrid approach does not necessarily help, it does not hurt either.

[39] As with the conventional POT technique, threshold selection for the hybrid model is an important factor. The same techniques for threshold choice apply here, but it should be kept in mind that the fitted gamma distribution is used to estimate the probability of exceeding the threshold [i.e., $1 - F(u)$]. In other words, the mass of the heavy-tailed GP part of the hybrid distribution is determined by the fitted gamma distribution, hence estimating this probability too low (i.e., choosing too high a threshold) implies less emphasis on a possibly heavy tail. Therefore the threshold u should be chosen in a region where the fit of the gamma distribution is not yet too bad, so that this probability is not unduly underestimated.

4.3. Mixture of Distributions

[40] Threshold selection for the stretched exponential distribution fit to daily precipitation intensity was found difficult in practice (section 3.2). An alternative approach would be to consider a mixture of distributions, with one of the components being the stretched exponential (i.e., with threshold $u = 0$). In this way, threshold selection could be avoided, while still taking advantage of the capability of the stretched exponential to produce an apparent heavy upper tail. Given the common use of a mixture of two exponentials (3) as a model for intensity, it would be natural to replace one of these exponentials with a stretched exponential (note that the stretched exponential (5) with $u = 0$ reduces to the exponential when the shape parameter $c = 1$); that is, a mixture whose cumulative distribution function is of the form

$$F(x;\omega,\sigma_1,c,\sigma_2) = 1-(1-\omega)\exp\left(-\frac{x}{\sigma_1}\right)-\omega\exp\left[-\left(\frac{x}{\sigma_2}\right)^c\right],$$
$$x>0, 0<\omega<1, \sigma_1, c, \sigma_2>0. \qquad (14)$$

It might also be plausible to constrain the shape parameter c in (14) to be equal to 2/3, the value hypothesized by *Wilson and Toumi* [2005], rather than estimating it from the data. Then only three parameters (i.e., two scale parameters, σ_1 and σ_2, and the mixing weight ω) would need to be

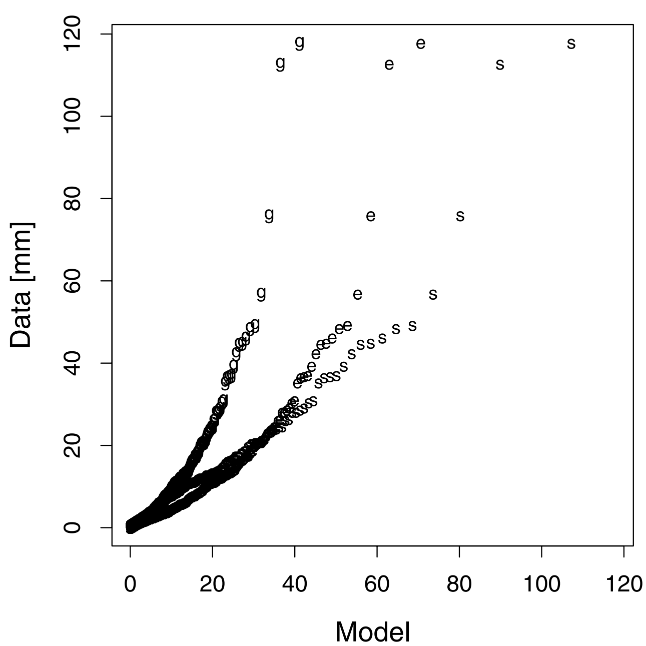

Figure 11. *Q-Q* plots of observed versus modeled gamma (*g*), mixed exponential (*e*) and mixed exponential/stretched exponential with c = 2/3 (*s*) quantiles of precipitation intensity at Fort Collins in July.

Table 3. Estimated Parameters and Log-Likelihood Values for a Mixture of Two Exponential Distributions and for a Mixture of an Exponential With a Stretched Exponential Distribution (With Shape Parameter $c = 2/3$) for Precipitation Intensity at Fort Collins in July

Mixture Model	Mixing Weight ω	Scale σ_1	Scale σ_2	Log-Likelihood
Two exponentials	0.68	1.72	11.07	−2053.35
Exponential, stretched exponential ($c = 2/3$)	0.60	1.96	6.86	−2062.08

estimated. Further, a stretched exponential component with $c \ll 1$ should dominate the mixture in the upper tail. Such a mixture would still be simpler than the more general dynamic mixture model that *Vrac and Naveau* [2007] proposed for intensity.

[41] For daily precipitation intensity at Fort Collins in July, Table 3 lists the parameter estimates for two mixture models, a mixture of two exponentials (3) and an exponential/stretched exponential mixture (14) with $c = 2/3$. See Figure 11 for the corresponding *Q-Q* plots. Compared to the gamma distribution the fit for the mixture of two exponentials is an improvement for high intensity, but still woefully inadequate. On the other hand, the exponential/stretched exponential mixture does not improve upon the mixed exponential in terms of overall fit (see log-likelihood values in Table 3). Yet it comes much closer to capturing the apparent heavy tail, while the fit to small to moderate values of intensity remains acceptable (Figure 11). If the shape parameter c of the stretched exponential in (14) were not constrained to equal 2/3, then the log-likelihood value improves substantially, but the shape parameter estimate is much higher than 2/3 (in fact greater than one). So the benefit of the penultimate approximation is completely lost, with the fit to the upper tail of intensity seriously deteriorating (*Q-Q* plot not shown).

[42] For other sites or even other months at Fort Collins, one should not necessarily expect to obtain a satisfactory fit to small to moderate intensities as well as high intensities using an exponential/stretched exponential mixture with the brute force constraint of $c = 2/3$. Nevertheless, the results for the single month of July at Fort Collins indicate the capability of the stretched exponential distribution to manufacture an apparent heavy upper tail. Unfortunately, the implementation of such a mixture model for precipitation intensity within a stochastic weather generator appears problematic, and the mixture approach therefore cannot be directly compared to the methods described in the previous two sections.

5. Discussion

[43] We have examined different ways of attempting to improve the simulation of daily extreme high precipitation by parametric stochastic weather generators. One approach, based on the stretched exponential distribution and justified by penultimate extreme value theory, appears difficult to implement in practice. An alternative approach, consisting of a hybrid distribution combining conventional modeling of precipitation intensity for low to moderate values with modeling of high intensity based on ultimate extreme value theory, appears more promising. This hybrid technique not only is capable of making the simulated distribution of precipitation intensity heavier, but still can allow for annual cycles and other covariates as needed in weather generators.

[44] By building on an existing GLM-based stochastic weather generator, the hybrid approach for modeling precipitation intensity is reasonably straightforward to implement. Nevertheless, it does require the estimation of at least one additional parameter, being difficult to retain the parsimonious nature of the original GLM-based weather generator. The hybrid approach only avoids additional complexity by constraining the relationship between high precipitation intensity and any covariates to be driven by the relationship in the existing weather generator (i.e., primarily by low to moderate intensity). Only the case of the generation of time series of daily weather at a single site has been treated, but the extension of the hybrid approach to multisite weather generators would appear feasible.

[45] Because temperature variables enter into parametric stochastic weather generators in a more complicated manner than does precipitation intensity, it might well be more difficult to implement improvements in the treatment of temperature extremes. Nevertheless, present weather generators do not necessarily simulate realistically extreme temperature events in the form of spells (e.g., heat waves or cold spells) [*Qian et al.*, 2008; *Semenov*, 2008]. Like the present paper, reliance on the statistical theory of extreme values would appear to be a promising avenue for devising any such improved methods.

[46] **Acknowledgments.** We thank Philippe Naveau for advice on this research. We would also like to thank three anonymous referees for valuable suggestions that refined and improved the article. Research was partially supported by NSF grant DMS-0355474 to the NCAR Geophysical Statistics Project, by NSF's Biocomplexity in the Environment program grants BE-0410348 and CNH-0709681, and by NCAR's Weather and Climate Assessment Science Program. The National Center for Atmospheric Research is managed by the University Corporation for Atmospheric Research under the sponsorship of the National Science Foundation. All calculations in this work have been done with the free software environment for statistical computing and graphics R, using the packages ismev and extRemes, see http://www.r-project.org/.

References

Abaurrea, J., and A. C. Cebrián (2002), Drought analysis based on a cluster Poisson model: Distribution of the most severe drought, *Clim. Res.*, *22*, 227–235.

Cameron, D., K. Bevin, and J. Tawn (2001), Modelling extreme rainfalls using a modified pulse Bartlett-Lewis stochastic rainfall model (with uncertainty), *Adv. Water Resour.*, *24*, 203–211.

Coles, S. (2001), *An Introduction to Statistical Modeling of Extreme Values*, Springer, London.

Cook, N. J., and R. Harris (2004), Exact and general FT1 penultimate distributions of extreme wind speeds drawn from tail-equivalent Weibull parents, *Struct. Saf.*, *26*, 391–420.

Efron, B., and R. J. Tibshirani (1993), *An Introduction to the Bootstrap*, Chapman and Hall, Boca Raton, Fla.

Embrechts, P., C. Klüppelberg, and T. Mikosch (1997), *Modelling Extremal Events for Insurance and Finance*, Springer-Verlag, Berlin.

Furrer, E. M., and R. W. Katz (2007), Generalized linear modeling approach to stochastic weather generators, *Clim. Res.*, *34*, 129–144.

Johnson, G. L., C. L. Hanson, S. P. Hardegree, and E. B. Ballard (1996), Stochastic weather simulation: Overview and analysis of two commonly used models, *J. Appl. Meteorol.*, *35*, 1878–1896.

Katz, R. W. (1996), Use of conditional stochastic models to generate climate change scenarios, *Clim. Change.*, *32*(3), 237–255.
Katz, R. W., M. B. Parlange, and P. Naveau (2002), Statistics of extremes in hydrology, *Adv. Water Resour.*, *25*, 1287–1304.
Koutsoyiannis, D. (2004), Statistics of extremes and estimation of extreme rainfall. 2: Empirical investigation of long rainfall records, *Hydrol. Sci. J.*, *49*, 591–610.
Leadbetter, M. R., G. Lindgren, and H. Rootzén (1983), *Extremes and Related Properties of Random Sequences and Processes*, Springer-Verlag, New York.
Madsen, H., P. Rasmussen, and D. Rosbjerg (1997), Comparison of annual maximum series and partial duration series methods for modeling extreme hydrologic events: 1. At-site modeling, *Water Resour. Res.*, *33*, 747–757.
McCullagh, P., and J. A. Nelder (1989), *Generalized Linear Models*, 2nd ed., Chapman and Hall, London.
Qian, B., S. Gameda, and H. Hayhoe (2008), Performance of stochastic weather generators LARS-WG and AAFC-WG for reproducing daily extremes of diverse Canadian climates, *Clim. Res.*, *37*, 17–33.
Rajagopalan, B., and U. Lall (1999), A k-nearest-neighbor simulator for daily precipitation and other weather variables, *Water Resour. Res.*, *35*, 3089–3101.
Reiss, R.-D., and M. Thomas (2007), *Statistical Analysis of Extreme Values with Applications to Insurance, Finance, Hydrology and Other Fields*, 3rd ed., Birkhäuser, Basel.
Richardson, C. W. (1981), Stochastic simulation of daily precipitation, temperature and solar radiation, *Water Resour. Res.*, *17*, 182–190.
Schwarz, G. (1978), Estimating the dimension of a model, *Ann. Stat.*, *6*, 461–464.
Semenov, M. A. (2008), Simulation of extreme weather events by a stochastic weather generator, *Clim. Res.*, *35*, 203–212.
Sharif, M., and D. H. Burn (2006), Simulating climate change scenarios using an improved k-nearest neighbor model, *J. Hydrol.*, *325*, 179–196.
Smith, R. L. (1989), Extreme value analysis of environmental time series: An example based on ozone data (with discussion), *Stat. Sci.*, *4*, 367–393.
Smith, R. E., and H. A. Schreiber (1974), Point process of seasonal thunderstorm rainfall: 2. Rainfall depth probabilities, *Water Resour. Res.*, *10*, 418–423.
Vrac, M., and P. Naveau (2007), Stochastic downscaling of precipitation: From dry events to heavy rainfalls, *Water Resour. Res.*, *43*, W07402, doi:10.1029/2006WR005308.
Wilby, R. L., T. M. L. Wigley, D. Conway, P. D. Jones, B. C. Hewitson, J. Main, and D. S. Wilks (1998), Statistical downscaling of general circulation model output: A comparison of methods, *Water Resour. Res.*, *34*, 2995–3008.
Wilks, D. S. (1989), Rainfall intensity, the Weibull distribution, and estimation of daily surface runoff, *J. Appl. Meteorol.*, *28*, 52–58.
Wilks, D. S. (1992), Adapting stochastic weather generation algorithms for climate change studies, *Clim. Change*, *22*, 67–84.
Wilks, D. S. (1999), Interannual variability and extreme-value characteristics of several stochastic daily precipitation models, *Agric. For. Meteorol.*, *93*, 153–169.
Wilks, D. S., and R. L. Wilby (1999), The weather generation game: A review of stochastic weather models, *Prog. Phys. Geogr.*, *23*(3), 329–357.
Wilson, P. S., and R. Toumi (2005), A fundamental probability distribution for heavy rainfall, *Geophys. Res. Lett.*, *32*, L14812, doi:10.1029/2005GL022465.
Wong, R. K. W. (1977), Weibull distribution, iterative likelihood technique and hydrometeorological data, *J. Appl. Meteorol.*, *16*, 1360–1364.
Yan, Z., S. Bate, R. E. Chandler, V. Isham, and H. Wheater (2006), Changes in extreme wind speeds in NW Europe simulated by generalized linear models, *Theor. Appl. Climatol.*, *83*, 121–137.

E. M. Furrer, Institute for Mathematics Applied to Geosciences, National Center for Atmospheric Research, P.O. Box 3000, Boulder, CO 80307, USA. (eva@ucar.edu)
R. W. Katz, Institute for Study of Society and Environment, National Center for Atmospheric Research, P.O. Box 3000, Boulder, CO 80307, USA. (rwk@ucar.edu)