

A Poisson process reparameterisation for Bayesian inference for extremes

Paul Sharkey and Jonathan A. Tawn
Email: p.sharkey1@lancs.ac.uk, j.tawn@lancs.ac.uk

STOR-i Centre for Doctoral Training, Department of Mathematics and Statistics, Lancaster University, Lancaster, LA1 4YF, U.K.

Abstract

A common approach to modelling extreme values is to consider the excesses above a high threshold as realisations of a non-homogeneous Poisson process. While this method offers the advantage of modelling using threshold-invariant extreme value parameters, the dependence between these parameters makes estimation more difficult. We present a novel approach for Bayesian estimation of the Poisson process model parameters by reparameterising in terms of a tuning parameter m . This paper presents a method for choosing the optimal value of m that near-orthogonalises the parameters, which is achieved by minimising the correlation between the asymptotic posterior distribution of the parameters. This choice of m ensures more rapid convergence and efficient sampling from the joint posterior distribution using Markov Chain Monte Carlo methods. Samples from the parameterisation of interest are then obtained by a simple transform. Results are presented in the cases of identically and non-identically distributed models for extreme rainfall in Cumbria, UK.

Keywords: Poisson processes, Extreme value theory, Bayesian inference, Reparameterisation, Covariate modelling.

1 A Poisson Process model for Extremes

The aim of extreme value analysis is to model rare occurrences of an observed process to extrapolate to give estimates of the probabilities of unobserved levels. In this way, one can make predictions of future extreme behaviour by estimating the behaviour of the process using an asymptotically justified limit model. Let X_1, X_2, \dots, X_n be a series of independent and identically distributed (iid) random variables with common distribution function F . Defining $M_n = \max\{X_1, X_2, \dots, X_n\}$, if there exists sequences of normalising constants $a_n > 0$ and b_n such that:

$$\Pr \left\{ \frac{M_n - b_n}{a_n} \leq x \right\} \rightarrow G(x) \quad \text{as } n \rightarrow \infty, \quad (1)$$

where G is non-degenerate, then G follows a generalised extreme value (GEV) distribution, with distribution function

$$G(x) = \exp \left\{ - \left[1 + \xi \left(\frac{x - \mu}{\sigma} \right) \right]_+^{-1/\xi} \right\}, \quad (2)$$

where $x_+ = \max(x, 0)$, $\sigma > 0$ and $\mu, \xi \in \mathbb{R}$. Here, μ, σ and ξ are location, scale and shape parameters respectively.

Using a series of block maxima from X_1, \dots, X_n , typically with blocks corresponding to years, the standard inference approach to give estimates of (μ, σ, ξ) is the maximum likelihood technique, which requires numerical optimisation methods. In these problems, particularly when covariates are involved, such methods may converge to local optima, with the consequence that parameter estimates are largely influenced by the choice of starting values. The standard asymptotic properties of the maximum likelihood estimators are subject to certain regularity conditions outlined in Smith (1985), but can give a poor representation of true uncertainty. In addition, flat likelihood surfaces can cause identifiability issues (Smith, 1987a). For these reasons, we choose to work in a Bayesian setting. Bayesian approaches have been used to make inferences about $\boldsymbol{\theta} = (\mu, \sigma, \xi)$ using standard Markov Chain Monte Carlo (MCMC) techniques. They have the advantage of being able to incorporate prior information when little is known about the extremes of interest, while also better accounting for parameter uncertainty when estimating functions of $\boldsymbol{\theta}$, such as return levels (Coles and Tawn, 1996). For a recent review, see Stephenson (2016).

An approach to inference that is considered to be more efficient than using block maxima is to consider a model for threshold excesses, which is superior in the sense that it reduces uncertainty due to utilising more extreme data (Smith, 1987b). Given a high threshold u , the conditional distribution of excesses above u can be approximated by a generalised Pareto (GP) distribution (Pickands, 1975) such that

$$\Pr(X - u > x | X > u) = \left(1 + \frac{\xi x}{\psi_u} \right)_+^{-1/\xi}, \quad x > 0,$$

where $\psi_u > 0$ and $\xi \in \mathbb{R}$ denote the scale and shape parameters respectively, with ψ_u dependent on the threshold u , while ξ is identical to the shape parameter of the GEV distribution. This model conditions on an exceedance, but a third parameter λ_u , denoting the rate of exceedance of X above the threshold u , must also be estimated.

Both of these extreme value approaches are special cases of a unifying limiting Poisson process characterisation of extremes (Smith, 1989; Coles, 2001). Let P_n be a sequence of point processes

such that

$$P_n = \left\{ \left(\frac{i}{n+1}, \frac{X_i - b_n}{a_n} \right) : i = 1, \dots, n \right\},$$

where $a_n > 0$ and b_n are the normalising constants in limit (1). The limit process is non-degenerate since the limit distribution of $(M_n - b_n)/a_n$ is non-degenerate. Small points are normalised to the same value $b_L = \lim_{n \rightarrow \infty} (x_L - b_n)/a_n$, where x_L is the lower endpoint of the distribution F . Large points are retained in the limit process. It follows that P_n converges to a non-homogeneous Poisson process P on regions of the form $A_y = (0, 1) \times [y, \infty)$, for $y > b_L$. The limit process P has an intensity measure on A_y given by

$$\Lambda(A_y) = \left[1 + \xi \left(\frac{y - \mu}{\sigma} \right) \right]_+^{-1/\xi}. \quad (3)$$

It is typical to assume that the limit process is a reasonable approximation to the behaviour of P_n , without normalisation of the $\{X_i\}$, on $A_u = (0, 1) \times [u, \infty)$, where u is a sufficiently high threshold and a_n, b_n are absorbed into the location and scale parameters of the intensity (3). It is often convenient to rescale the intensity by a factor m , where $m > 0$ is free, so that the n observations consist of m blocks of size n/m with the maximum M_m of each block following a $\text{GEV}(\mu_m, \sigma_m, \xi)$ distribution, with ξ invariant to the choice of m . The Poisson process likelihood can be expressed as

$$L(\boldsymbol{\theta}_m) = \exp \left\{ -m \left[1 + \xi \left(\frac{u - \mu_m}{\sigma_m} \right) \right]_+^{-1/\xi} \right\} \prod_{j=1}^r \frac{1}{\sigma_m} \left[1 + \xi \left(\frac{x_j - \mu_m}{\sigma_m} \right) \right]_+^{-1/\xi - 1}, \quad (4)$$

where $\boldsymbol{\theta}_m = (\mu_m, \sigma_m, \xi)$ denotes the rescaled parameters, r denotes the number of excesses above the threshold u and $x_j > u, j = 1, \dots, r$, denote the exceedances. It is possible to move between parameterisations associated with different numbers of blocks. If for k blocks the block maximum is denoted by M_k and follows a GEV distribution with the parameters $\boldsymbol{\theta}_k = (\mu_k, \sigma_k, \xi)$, then for all x

$$\Pr(M_k < x) = \Pr(M_m < x)^{k/m}.$$

As M_k is $\text{GEV}(\mu_k, \sigma_k, \xi)$ and M_m is $\text{GEV}(\mu_m, \sigma_m, \xi)$ it follows that

$$\begin{aligned} \mu_k &= \mu_m - \frac{\sigma_m}{\xi} \left(1 - \left(\frac{k}{m} \right)^{-\xi} \right) \\ \sigma_k &= \sigma_m \left(\frac{k}{m} \right)^{-\xi}. \end{aligned} \quad (5)$$

In this paper, we present a method to improve inference for $\boldsymbol{\theta}_k$, the parameterisation of interest. For an ‘optimal’ choice of m we first undertake inference for $\boldsymbol{\theta}_m$ before transforming our results to give inference for $\boldsymbol{\theta}_k$ using the mapping in expression (5).

In many practical problems, k is taken to be n_y , the number of years of observation, so that the annual maximum has a GEV distribution with parameters $\boldsymbol{\theta}_{n_y} = (\mu_{n_y}, \sigma_{n_y}, \xi)$. Although inference is for the annual maximum distribution parameters $\boldsymbol{\theta}_{n_y}$, the Poisson process model makes use of all data that are extreme, so inferences are more precise than estimates based on a direct fit of the GEV distribution to the annual maximum data as noted above.

To help see how the choice of m affects inference, consider the case when $m = r$, the number of excesses above the threshold u . If a likelihood inference was being used with this choice of m , the maximum likelihood estimators $(\hat{\mu}_r, \hat{\sigma}_r, \hat{\xi}) = (u, \hat{\psi}_u, \hat{\xi})$, see Appendix A for more details. Therefore, Bayesian inference for the parameterisation of the Poisson process model when $m = r$ is equivalent to Bayesian inference for the GP model.

Although inference for the Poisson process and GP models is essentially the same approach when $m = r$, they differ in parameterisation, and hence inference, when $m \neq r$. The GP model is advantageous in that λ_u is globally orthogonal to ψ_u and ξ . Chavez-Demoulin and Davison (2005) achieved local orthogonalisation of the GP model at the maximum likelihood estimates by reparameterising the scale parameter as $\nu_u = \psi_u(1 + \xi)$. This ensures all the GP tail model parameters are orthogonal locally at the likelihood mode. However, the scale parameter is still dependent on the choice of threshold. Unlike the GP, the parameters of the Poisson process model are invariant to choice of threshold, which makes it more suitable for covariate modelling and hence suggests that it may be the better parameterisation to use. In contrast, it has been found that the parameters are highly dependent, making estimation more difficult.

As we are working in the Bayesian framework, strongly dependent parameters lead to poor mixing in our MCMC procedure (Hills and Smith, 1992). A common way of overcoming this is to explore the parameter space using a dependent proposal random walk Metropolis-Hastings algorithm, though this requires a knowledge of the parameter dependence structure *a priori*. Even in this case, the dependence structure potentially varies in different regions of the parameter space, which may require different parameterisations of the proposal to be applied. The alternative approach is to consider a reparameterisation to give orthogonal parameters. However, Cox and Reid (1987) show that global orthogonalisation cannot be achieved in general.

This paper illustrates an approach to improving Bayesian inference and efficiency for the Poisson process model. Our method exploits the scaling factor m as a means of creating a near-orthogonal

representation of the parameter space. While it is not possible in our case to find a value of m that diagonalises the Fisher information matrix, we focus on minimising the off-diagonal components of the covariance matrix. We present a method for choosing the ‘best’ value of m such that near-orthogonality of the model parameters is achieved, and thus improves the convergence of MCMC and sampling from the joint posterior distribution. Our focus is on Bayesian inference but the reparameterisations we find can be used to improve likelihood inference as well, simply by ignoring the prior term.

The structure of the paper is as follows. Section 2 examines the idea of reparameterising in terms of the scaling factor m and how this can be implemented in a Bayesian framework. Section 3 discusses the choice of m to optimise the sampling from the joint posterior distribution in the case where X_1, \dots, X_n are iid. Section 4 explores this choice when allowing for non-identically distributed variables through covariates in the model parameters. Section 5 describes an application of our methodology to extreme rainfall in Cumbria, UK, which experienced major flooding events in November 2009 and December 2015.

2 Bayesian Inference

Bayesian estimation of the Poisson process model parameters involves the specification of a prior distribution $\pi(\boldsymbol{\theta}_m)$. Then using Bayes Theorem, the posterior distribution of $\boldsymbol{\theta}_m$ can be expressed as

$$\pi(\boldsymbol{\theta}_m|\mathbf{x}) \propto \pi(\boldsymbol{\theta}_m)L(\boldsymbol{\theta}_m),$$

where $L(\boldsymbol{\theta}_m)$ is the likelihood as defined in (4) and \mathbf{x} denotes the excesses of the threshold u . We sample from the posterior distribution using a random walk Metropolis-Hastings scheme. Proposal values of each parameter are drawn sequentially from a univariate Normal distribution and accepted with a probability defined as the posterior ratio of the proposed state relative to the current state of the Markov chain. In all cases throughout the paper, each individual parameter chain is tuned to give the acceptance rate in the range of 20% – 25% to satisfy the optimality criterion of Roberts et al. (2001). For illustration purposes, results in Sections 2 and 3 are from the analysis of simulated iid data. A total of 300 exceedances above a threshold $u = 30$ are simulated from a Poisson process model with $\boldsymbol{\theta}_1 = (80, 15, 0.05)$. Figure 1 shows individual parameter chains for $\boldsymbol{\theta}_k$ from a random walk Metropolis scheme run for 50000 iterations with a burn-in of 1000 removed, where $k = 1$ and a chosen $m = 1$. This figure shows the clear poor mixing of each component of $\boldsymbol{\theta}_1$, indicating non-convergence and strong dependence in the posterior sampling.

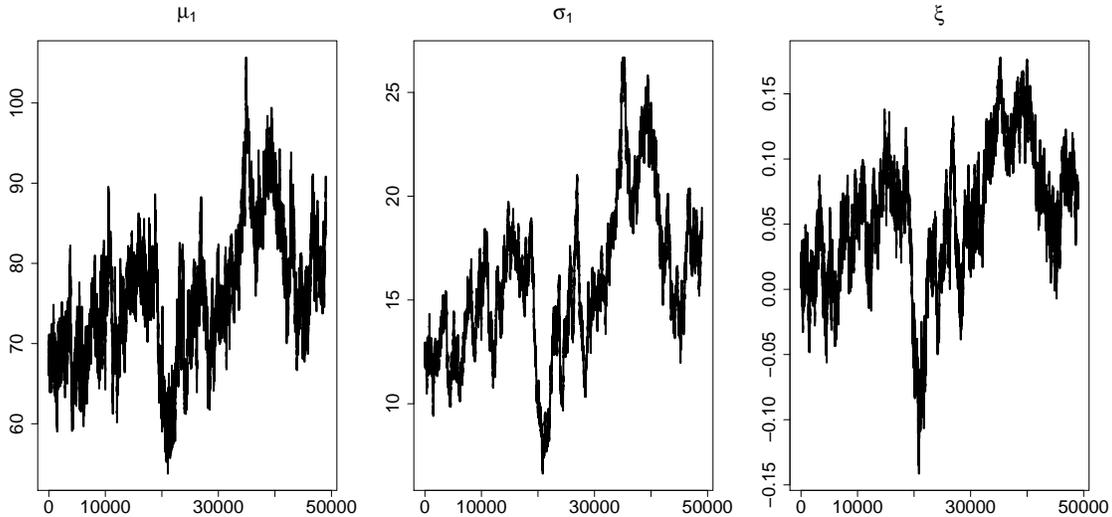


Figure 1: Random-walk Metropolis chains run for each component of θ_1 .

We explore how reparameterising the model in terms of m can improve sampling performance. For a general prior on the parameterisation of interest θ_k , denoted by $\pi(\theta_k)$, Appendix B derives that the prior on the transformed parameter space θ_m is

$$\pi(\theta_m) = \left(\frac{m}{k}\right)^{-\xi} \pi(\theta_k). \quad (6)$$

In this example, independent Uniform priors are placed on μ_1 , $\log \sigma_1$ and ξ , which gives

$$\pi(\theta_1) \propto \frac{1}{\sigma_1}; \quad \mu_1 \in \mathbb{R}, \sigma_1 > 0, \xi \in \mathbb{R}. \quad (7)$$

This choice of prior results in a proper posterior distribution, provided there are at least 4 threshold excesses (Northrop and Attalides, 2016). By finding a value of m that near-orthogonalises the parameters of the posterior distribution $\pi(\theta_m|\mathbf{x})$, we can run an efficient MCMC scheme on θ_m before transforming the samples to θ_k . It is noted in Wadsworth et al. (2010) that setting m to be the number of exceedances above the threshold, i.e. $m = r$, improves the mixing properties of the chain, as is illustrated in Figure 2. This is approximately equivalent to inference using a GP model, as discussed in Section 1.

Given this choice of m , the MCMC scheme is run for θ_m before transforming to estimate the posterior of θ_1 using the mapping in (5), where $k = 1$ in this case. Figure 3 shows contour plots of estimated joint posterior densities of θ_1 . It compares the samples from directly estimating the posterior of θ_1 with that from transforming from the MCMC samples of the posterior of θ_m to give

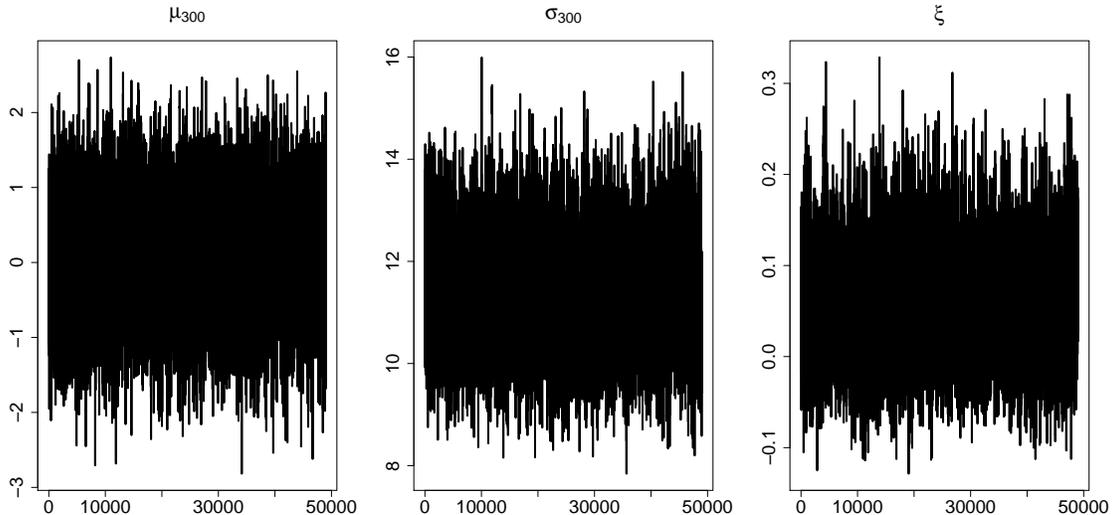


Figure 2: Random-walk Metropolis chains run for parameters θ_r , where $r = 300$ is the number of exceedances in the simulated data.

a posterior sample for θ_1 . Figure 3 indicates that θ_1 are highly correlated, with the result that we only sample from a small proportion of the parameter space when exploring using independent random walks for each parameter. This explains the poor mixing if we were to run the MCMC without a transformation. It is clear that, after back-transforming to θ_1 , the reparameterisation enables a more thorough exploration of the parameter space. However, while this transformation is a useful tool in enabling an efficient Bayesian inference procedure, further investigation is necessary in the choice of m to achieve near-orthogonality of the parameter space and thus maximising the efficiency of the MCMC procedure.

3 Choosing m optimally

As illustrated in Section 2, the choice of m in the Poisson process likelihood can improve the performance of the MCMC required to estimate the posterior density of model parameters θ_k . We desire a value of m such that near-orthogonality of θ_m is achieved, before using the expressions in (5) to transform to the parameterisation of interest, e.g. θ_1 or θ_{n_y} . As a measure of dependence, we use the asymptotic expected correlation matrix of the posterior distribution of $\theta_m|\mathbf{x}$. In particular, we explore how the off-diagonal components of the matrix, that is, the correlation between parameters, changes with m . The covariance matrix associated with $\theta_m|\mathbf{x}$ can be derived analytically by inverting the Fisher information matrix of the Poisson process log-likelihood (see Appendix C). The correlation matrix is then obtained by normalising so that the matrix has a unit diagonal.

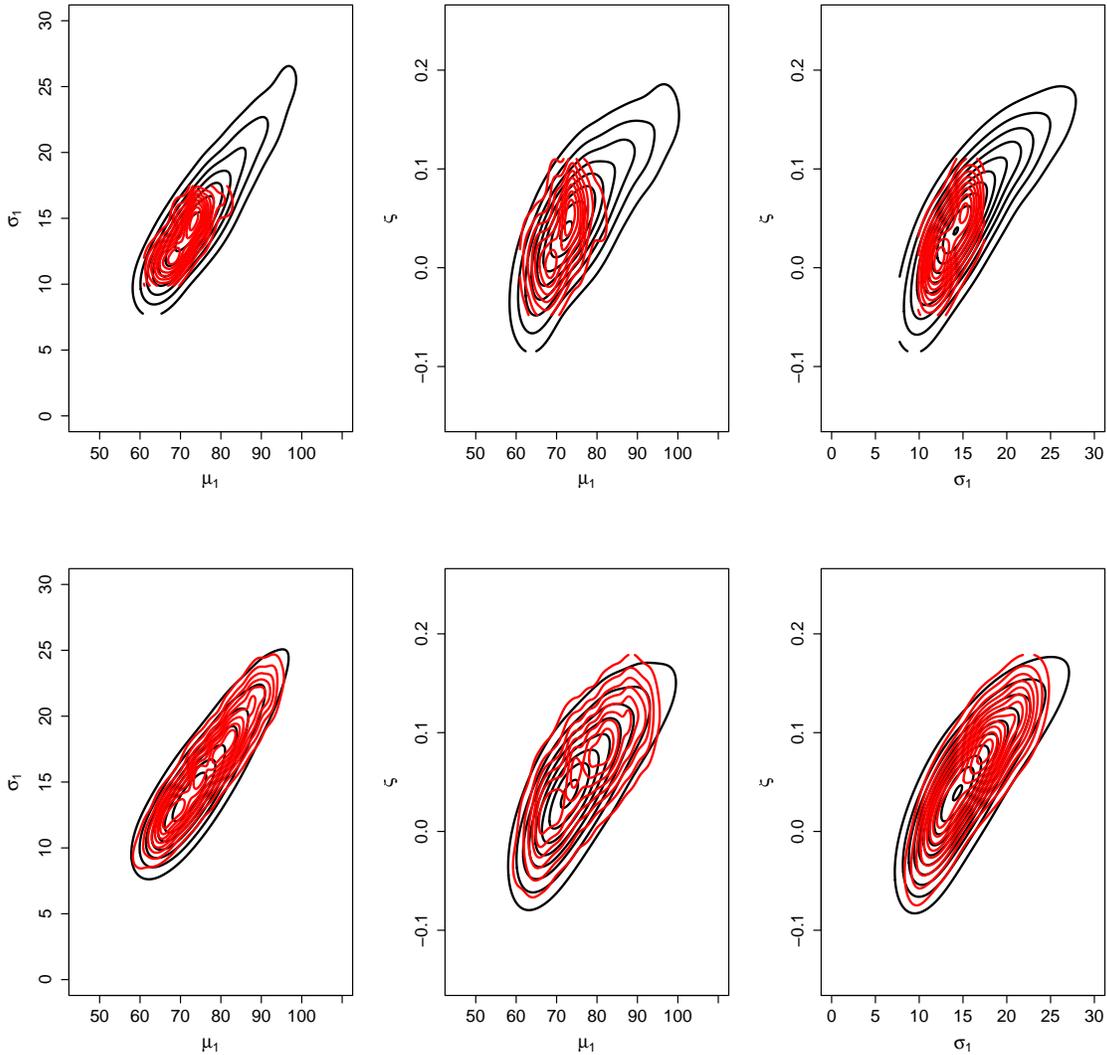


Figure 3: Contour plots of the estimated joint posterior of θ_1 for 4000 iterations (top) and 49000 iterations (bottom) created from the transformed samples drawn from the MCMC procedure for θ_m (in black) and samples of θ_1 drawn directly (in red).

Other choices for the measure of the dependence of the posterior could have been used, such as the inverse of the Hessian matrix (or the expected Hessian matrix) of the log-posterior, evaluated at the posterior mode. For inference problems with strong information from the data relative to the prior there will be limited differences in the approach and similar values for the optimal m will be found. In contrast, if the prior is strongly informative and the number of threshold exceedances is small then the choice of m from using our approach could be far from optimal. Also the use of the

observed, rather than expected, Hessian may better represent the actual posterior distribution of θ_m and deliver a choice of m that better achieves orthogonalisation, see Efron and Hinkley (1978) and Tawn (1987) respectively.

We prefer our choice of measure of dependence as for iid problems it gives closed form results for m which can be used without the computational work required for other approaches, and this gives valuable insight into the choice of m . Furthermore, informative priors rarely arise in extreme value problems, and so information in the data typically dominates information in the prior, particularly around the posterior mode. It should be pointed out however, that the prior is used in the MCMC so there is no loss of prior information in our approach. Also standard MCMC diagnostics should be used even after the selection of an optimal m , so if the asymptotic posterior correlations differ much from the posterior correlations, making our choice of m poor, this will be obvious and a more complete but computationally burdensome analysis can be conducted using the methods described above.

In this section, we use the data introduced in Section 2. For all integers $m \in [1, 500]$, maximum likelihood estimates $\hat{\theta}_m$ are computed and pairwise correlations calculated by substituting $\hat{\theta}_m$ into the expressions for the Fisher information matrix, in Appendix C, and taking the inverse. Figure 4 shows how parameter correlations change with the choice of m , illustrating that the asymptotic posterior distributions of μ_m and ξ are orthogonal when $m = r$, the number of excesses above a threshold, which explains the findings of Wadsworth et al. (2010).

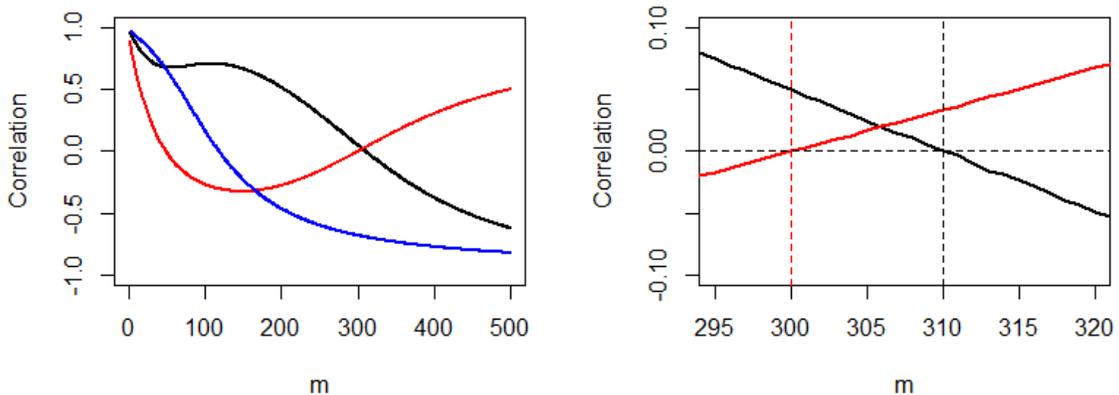


Figure 4: Left: Estimated parameter correlations changing with m : ρ_{μ_m, σ_m} (black), $\rho_{\mu_m, \xi}$ (red), $\rho_{\sigma_m, \xi}$ (blue). Right: Expanded region of the graph showing $\rho_{\mu_m, \xi} = 0$ for m close to r where $r = 300$ is the number of excesses above the threshold, while $\rho_{\mu_m, \sigma_m} = 0$ when $m \approx 310$.

It is proposed that MCMC mixing can be further improved by minimising the overall correlation in the asymptotic posterior distribution of θ_m . Therefore, we would like to find the value of m such that $\rho(\theta_m)$ is minimised, where $\rho(\theta_m)$ is defined as

$$\rho(\theta_m) = |\rho_{\mu_m, \sigma_m}| + |\rho_{\mu_m, \xi}| + |\rho_{\sigma_m, \xi}|, \quad (8)$$

where ρ_{μ_m, σ_m} denotes the asymptotic posterior correlation between μ_m and σ_m for example. We also

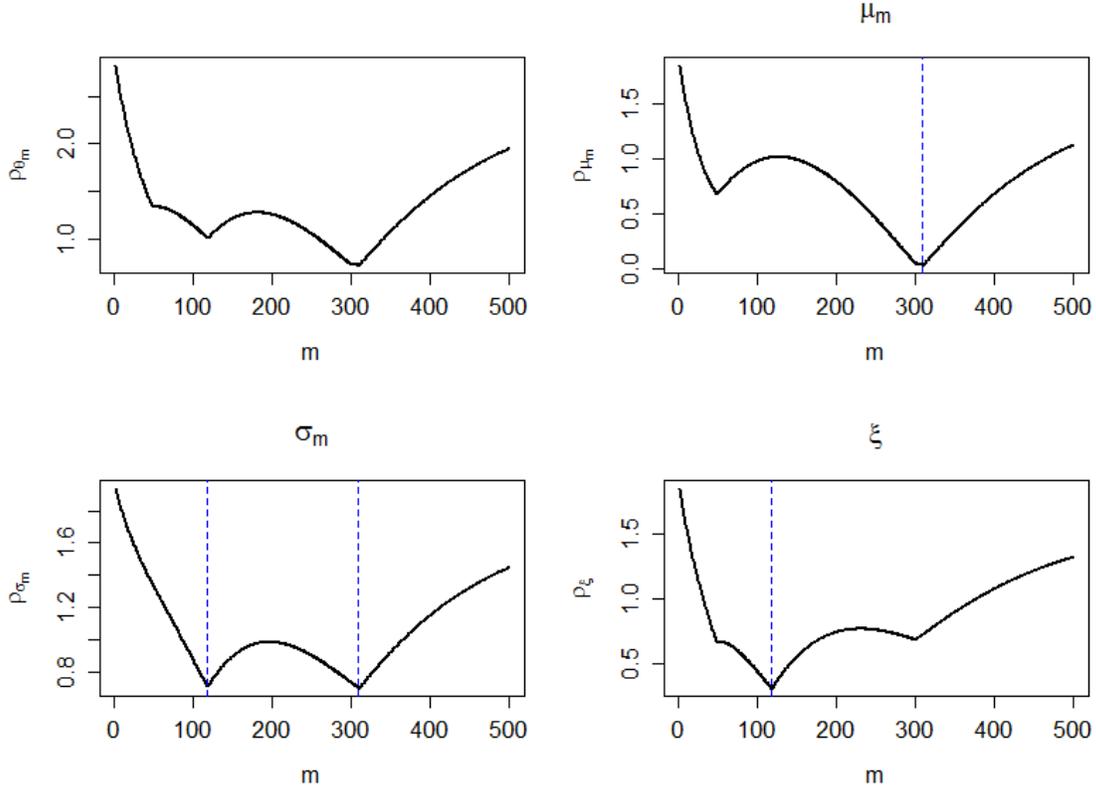


Figure 5: How $\rho(\theta_m)$ changes with m (top left) and how correlations in each individual estimated parameter, as measured by ρ_{μ_m} , ρ_{σ_m} and ρ_{ξ} , change with m .

look at the sum of the asymptotic posterior correlation terms involving each individual parameter estimate. For example, we define ρ_{μ_m} , the asymptotic posterior correlation associated with the estimate of μ_m , to be:

$$\rho_{\mu_m} = |\rho_{\mu_m, \sigma_m}| + |\rho_{\mu_m, \xi}|. \quad (9)$$

Figure 5 shows how the asymptotic posterior correlation associated with each parameter varies with m . From Figure 5 we see that while ρ_{μ_m} is minimised at the value of m for which $\rho_{\mu_m, \sigma_m} = 0$ (see Figure 4), ρ_{σ_m} and ρ_{ξ} have minima at the value of m for which $\rho_{\sigma_m, \xi} = 0$. We denote the latter

minimum by m_1 and the former by m_2 . In terms of the covariance function, this can be written as:

$$\text{ACov}(\sigma_{m_1}, \xi | \mathbf{x}) = \text{ACov}(\mu_{m_2}, \sigma_{m_2} | \mathbf{x}) = 0, \quad (10)$$

where ACov denotes the asymptotic covariance. Figure 5 shows that m_2 also minimises the total asymptotic posterior correlation in the model.

One would expect that the values of m for which $\rho(\boldsymbol{\theta}_m)$ is minimised would correspond to the MCMC chain of $\boldsymbol{\theta}_m$ with good mixing properties. We examine the effective sample size (ESS) as a way of evaluating this objectively. ESS is a measure of the equivalent number of independent iterations that the chain represents (Robert and Casella, 2009). MCMC samples are often positively autocorrelated, and thus are less precise in representing the posterior than if the chain was independent. The ESS of a parameter chain ϕ is defined as

$$\text{ESS}_\phi = \frac{n}{1 + 2 \sum_{i=1}^{\infty} \nu_i}, \quad (11)$$

where n is the length of the chain and ν_i denotes the autocorrelation in the sampled chain of ϕ at lag i . In practice, the sum of the autocorrelations is truncated when ν_i drops beneath a certain level. Figure 6 shows how ESS varies with m for each parameter in $\boldsymbol{\theta}_m$. For these data the ESS follow a pattern we found to typically occur. We see that ESS_{μ_m} is maximised at $m = m_2$ due to the near-orthogonality of μ_{m_2} with σ_{m_2} and ξ . We find that ESS_{σ_m} is maximised for $m_1 < m < m_2$, as σ_{m_1} remains substantially positively correlated with μ_{m_1} and σ_{m_2} is negatively correlated with ξ . Similarly, ESS_ξ is maximised at a value of m close to m_1 , but ξ is negatively correlated with μ_{m_1} , which explains the slight distortion. From these results, we postulate that a selection of m in the interval (m_1, m_2) would ensure the most rapid convergence of the MCMC chain of $\boldsymbol{\theta}_m$, thus enabling an effective sampling procedure from the joint posterior. Figure 6 shows clearly the benefits of the proposed approach. For example, $\text{ESS}_{\mu_{310}} = 7459.17$ and $\text{ESS}_{\mu_1} = 24.44$, illustrating that the former parameterisation is over 310 times more efficient than the latter. In addition, by introducing the interval (m_1, m_2) , this approach gives a degree of flexibility to the choice of m and giving a balance of mixing quality across the model parameters.

The quantities m_1 and m_2 can be found by numerical solution of the equations $\text{ACov}(\sigma_m, \xi | \mathbf{x}) = 0$ and $\text{ACov}(\mu_m, \sigma_m | \mathbf{x}) = 0$ respectively, using the asymptotic covariance matrix of the posterior of $\boldsymbol{\theta}_m$, which is given by the inverse of the Fisher information (see Appendix C). Approximate analytical expressions for m_1 and m_2 can be derived using Halley's method for root-finding (Gander,

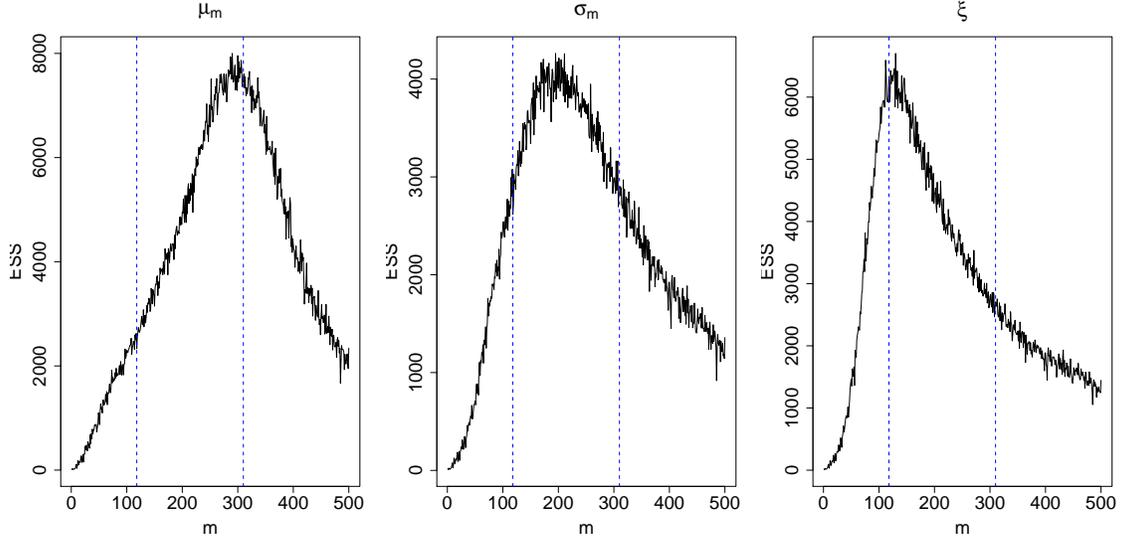


Figure 6: How ESS varies with m for each parameter in θ_m . The blue dashed lines represent $m = m_1$ (left) and $m = m_2$ (right) in the simulated data example, where m_1 and m_2 are defined by property (10). In the calculations, the sum of the autocorrelations were truncated when the autocorrelations in the chain drop below 0.05.

1985) applied to equations (10). This method yields the following approximations of m_1 and m_2 :

$$\hat{m}_1 = r \frac{(2\xi + 1) \left(1 + 2\xi + (\xi + 1) \log \left[\frac{2\xi+3}{2\xi+1} \right] \right)}{(2\xi + 1) \left(3 + 2\xi - (\xi + 1) \log \left[\frac{2\xi+3}{2\xi+1} \right] \right)} \quad (12)$$

$$\hat{m}_2 = r \frac{2\xi^2 + 13\xi + 8}{2\xi^2 + 9\xi + 8}. \quad (13)$$

In practice, the values of \hat{m}_1 and \hat{m}_2 are estimated by using an estimate of ξ , such as the maximum likelihood estimate. Figure 7 shows how \hat{m}_1 and \hat{m}_2 change relative to r for a range of ξ . This illustrates that for negative estimates of the shape parameter, r is not a suitable candidate to be the ‘optimal’ value of m as it is not in the range (m_1, m_2) . In the simulated data used in this section, although a selection of $m = r$ is reasonable, Figure 6 shows that this may not be wise if one was primarily concerned about sampling well from ξ , for example. In this case, \hat{m}_2 is relatively close to r , but Figure 7 shows that this is not the case for models with a larger positive estimate of ξ .

A simulation study was carried out to assess the suitability of expressions \hat{m}_1 and \hat{m}_2 as approximations to m_1 and m_2 respectively. A total of 1000 Poisson processes were simulated with different values of θ_m . The approximations were calculated and compared with the true values of m_1 and m_2 , which were obtained exactly by numerical methods. It was found that $|\hat{m}_i - m_i| < 0.1$ for

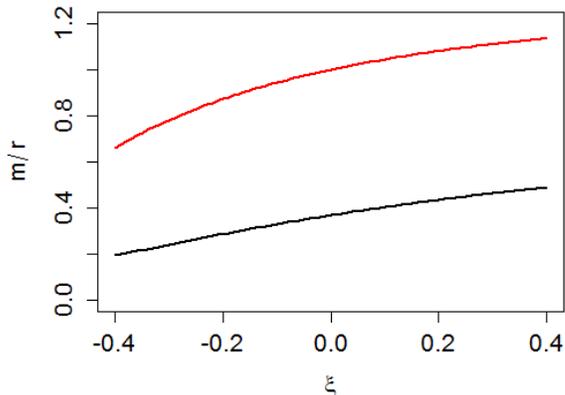


Figure 7: How \hat{m}_1 and \hat{m}_2 change as a multiple of r with respect to $\hat{\xi}$: \hat{m}_1/r (black), \hat{m}_2/r (red).

$i = 1, 2$ always, while $|\hat{m}_i - m_i| < 0.01$ for 78% and 88.2% of the time for $i = 1, 2$ respectively. Both quantities were compared to the performance of other approximations derived using Newton’s method, which unlike Halley’s method does not account for the curvature in a function. Simulations show that the root mean square errors are significantly smaller for estimates of m_i using Halley’s method (0.2% and 5% smaller than Newton’s method for $i = 1, 2$ respectively). A summary of the reparameterisation method is given in Algorithm 1.

Algorithm 1: Sampling from the posterior distribution of the Poisson process model parameters after reparameterising

Data: Threshold excesses \mathbf{x}

Result: Samples from the posterior distribution $\pi(\boldsymbol{\theta}_k|\mathbf{x})$

- 1 Choose parameterisation of interest $\boldsymbol{\theta}_k$;
 - 2 **if** *iid* **then**
 - 3 Obtain estimate of shape parameter ξ using maximum likelihood, for example;
 - 4 Compute \hat{m}_1 and \hat{m}_2 as defined in (12) and (13);
 - 5 Choose m in range (\hat{m}_1, \hat{m}_2) ;
 - 6 **else**
 - 7 Choose m to be the value of m that numerically solves $\rho_{\mu_m^{(0)}, \sigma_m} = 0$;
 - 8 Obtain MCMC samples for posterior distribution $\pi(\boldsymbol{\theta}_m|\mathbf{x})$;
 - 9 Transform to obtain samples from $\pi(\boldsymbol{\theta}_k|\mathbf{x})$
-

4 Choosing m in the presence of non-stationarity

In many practical applications, processes exhibit trends or seasonal effects caused by underlying mechanisms. The standard methods for modelling extremes of non-identically distributed random variables were introduced by Davison and Smith (1990) and Smith (1989), using a Poisson process and Generalised Pareto distribution respectively. Both approaches involve setting a constant threshold and modelling the parameters as functions of covariates. In this way, we model the non-stationarity through the conditional distribution of the process on the covariates. We follow the Poisson process model of Smith (1989) as the parameters are invariant to the choice of threshold if the model is appropriate. We define the covariate-dependent parameters $\boldsymbol{\theta}_m(z) = (\mu_m(z), \sigma_m(z), \xi(z))$, for covariates z . Often in practice, the shape parameter ξ is assumed to be constant. A log-link is typically used to ensure positivity of $\sigma_m(z)$.

The process of choosing m is complicated when modelling in the presence of covariates. This is partially caused by a modification of the integrated intensity measure, which becomes

$$\Lambda(A) = m \int_{\mathbf{z}} \left[1 + \xi(z) \left(\frac{u - \mu_m(z)}{\sigma_m(z)} \right) \right]^{-1/\xi(z)} g(z) dz, \quad (14)$$

where g denotes the probability density function of the covariates, which is unknown and with covariate space \mathbf{z} . The density term g is required as the covariates associated with exceedances of the threshold u are random. In addition, the extra parameters introduced by modelling covariates increases the overall correlation in the model parameters.

For simplicity, we restrict our attention to the case of modelling when the location parameter is a linear function of a covariate, that is,

$$\mu_m(z) = \mu_m^{(0)} + \mu_m^{(1)} z, \quad \sigma_m(z) = \sigma_m, \quad \xi(z) = \xi,$$

where we centre the covariate z , as this leads to parameters $\mu_m^{(0)}$ and $\mu_m^{(1)}$ being orthogonal. Note that the regression parameter $\mu_m^{(1)}$ is invariant to the choice of m . A total of 233 excesses above a threshold of $u = 15$ are simulated from a Poisson process model with $\mu_1^{(0)} = 75$, $\mu_1^{(1)} = 30$, $\sigma_1 = 15$, $\xi = -0.05$. We choose g to follow an Exp(2) distribution, noting that one could also choose g to be the density of a covariate that is used in practice. We impose an improper Uniform prior on the regression parameter $\mu_1^{(1)}$ and set up the MCMC scheme in the same manner as in Section 3.

The objective remains to identify the value of m that achieves near-orthogonality of the parameters of the posterior distribution. Like before, we run an MCMC sampler on $\boldsymbol{\theta}_m(z)$ and transform the

samples back to the parameterisation of interest $\theta_k(z)$, which can be obtained as in (5) using the relations

$$\begin{aligned}\mu_k^{(0)} &= \mu_m^{(0)} - \frac{\sigma_m}{\xi} \left(1 - \left(\frac{k}{m} \right)^{-\xi} \right) \\ \mu_k^{(1)} &= \mu_m^{(1)} \\ \sigma_k &= \sigma_m \left(\frac{k}{m} \right)^{-\xi}.\end{aligned}\tag{15}$$

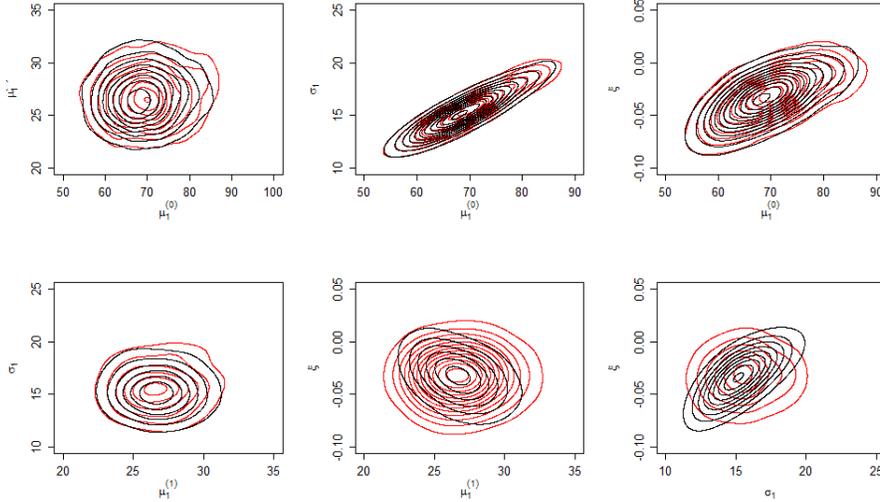


Figure 8: Contour plots of estimated posterior densities of $\theta_1(z)$ having sampled from the joint posterior directly (red) and having transformed using (15) after reparameterising from $\theta_{85}(z)$ (black).

The complication of the integral term in the likelihood for non-identically distributed variables means that it is no longer feasible to gain an analytical approximation for the optimal value of m . A referee has suggested a possible route to obtaining such expressions for m in the non-stationary case, is by building on results in Attalides (2015) and using a non-constant threshold as in Northrop and Jonathan (2011), but as this moves away from our constant threshold case we do not pursue this. We therefore choose a value of m that minimises the asymptotic posterior correlation in the model. Analogous to Section 3, the optimal m coincides with the value of m such that $\rho_{\mu_m^{(0)}, \sigma_m}^{(0)} = 0$. Using numerical methods, we identify that this corresponds to a value of $m = 85$ for the simulated data example. Figure 8 shows contour plots of estimated posterior densities of $\theta_1(z)$, comparing the sampling from directly estimating the posterior $\theta_1(z)$ with that from transforming the samples from the estimated posterior of $\theta_m(z)$ to give a sample from the posterior of $\theta_1(z)$. From this figure, we see that the reparameterisation improves the sampling from the posterior $\theta_1(z)$.

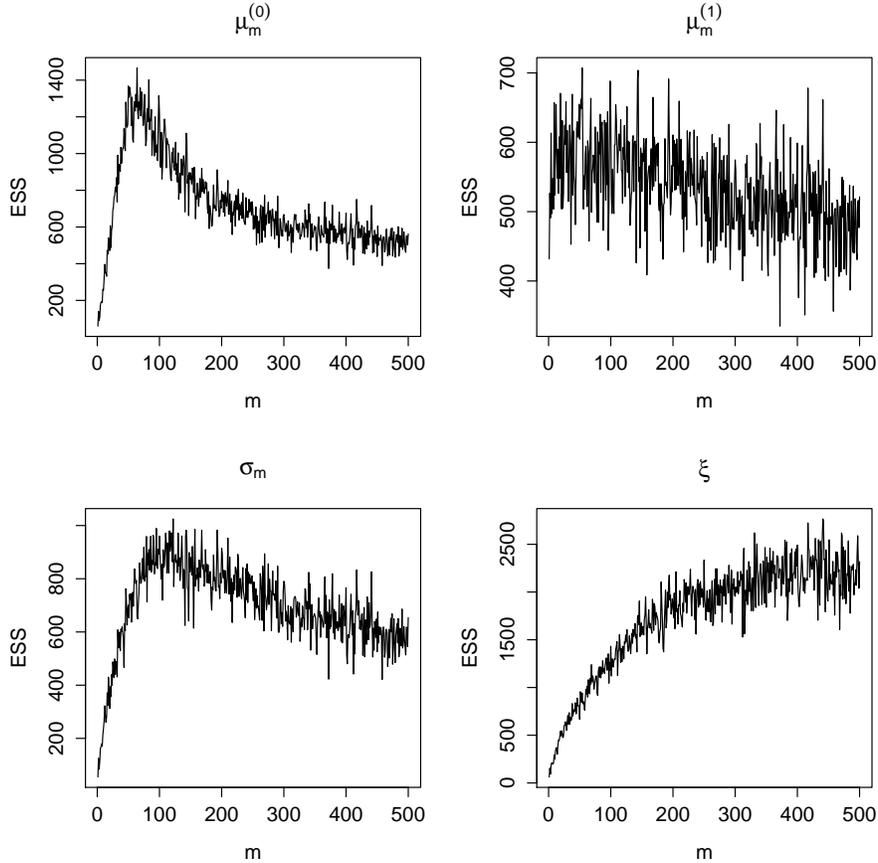


Figure 9: Effective sample size of each parameter chain of the MCMC procedure.

We again inspect the effective sample size for each parameter as a way of comparing the efficiency of the MCMC under different parameterisations. Figure 9 shows how the effective sample size varies with m for each parameter. This figure shows how the quality of mixing is approximately maximised in $\mu_m^{(0)}$ for the value of m that minimises the asymptotic posterior correlation. Mixing for $\mu_m^{(1)}$ is consistent across all values of m . Interestingly, mixing in ξ increases as the value of m increases. Without a formal measure for the quality of mixing across the parameters, it is found that, when averaging the effective sample size over the number of parameters, the ESS is stable with respect to m in the interval spanning from the value of m such that $\rho_{\mu_m^{(0)}, \sigma_m^{(0)}} = 0$ and the value of m such that $\rho_{\sigma_m^{(0)}, \xi} = 0$, like in Section 3. For a summary of how the reparameterisation method can be used in the presence of non-stationarity, see Algorithm 1.

5 Case study: Cumbria rainfall

In this section, we present a study as an example of how this reparameterisation method can be used in practice. In particular, we analyse data taken from the Met Office UKCP09 project, which contains daily baseline averages of surface rainfall observations, measured in millimetres, in $25\text{km} \times 25\text{km}$ grid cells across the United Kingdom in the period 1958-2012. In this analysis, we focus on a grid cell in Cumbria, which has been affected by numerous flood events in recent years, most notably in 2007, 2009 and 2015. In particular, the December 2015 event resulted in an estimated £5 billion worth of damage, with rain gauges reaching unprecedented levels. Many explanations have been postulated for the seemingly increased rate of flooding in the North West of England, including climate change, natural climate variability or a combination of both. The baseline average data for the flood events in December 2015 are not yet available, but this event is widely regarded as being more extreme than the event in November 2009, the levels of which were reported at the time to correspond to return periods of greater than 100 years. We focus our analysis on the 2009 event, looking in particular at how a phase of climate variability, in the form of the North Atlantic Oscillation (NAO) index, can have a significant impact on the probability of an extreme event occurring in any given year.

Rainfall datasets on a daily scale are commonly known to exhibit a degree of serial correlation. Analysis of autocorrelation and partial autocorrelation plots indicates that rainfall on day t is dependent on the rainfall of the previous five days. In addition, the data may exhibit seasonal effects. However, while serial dependence affects the effective sample size of a dataset, it does not affect correlations between parameters, and is thus unlikely to influence the choice of m . For the purposes of illustrating our method, we make the assumption for now that the rainfall observations are iid and proceed with the method outlined in Section 3. We wish to obtain information about the parameters corresponding to the distribution of annual maxima, i.e. θ_{55} . Standard threshold diagnostics (Coles, 2001) indicate a threshold of $u = 15$ is appropriate, which corresponds to the 95.6% quantile of the data. There are $r = 880$ excesses above u (see Figure 10). We obtain bounds m_1 and m_2 , then choose a value $m_1 < m < m_2$ that will achieve near-orthogonality of the Poisson process model parameters to improve MCMC sampling from the joint posterior distribution. We obtain $\hat{\xi} = 0.087$ using maximum likelihood when $m = r$, which we use to obtain approximations for m_1 and m_2 as in (12) and (13). From this, we obtain $\hat{m}_1 \approx 351$ and $\hat{m}_2 \approx 915$. We checked that \hat{m}_1 and \hat{m}_2 represent good approximations by solving equations (10) to obtain $m_1 = 350.82$ and $m_2 = 914.96$. Since $r = 880$ is contained in the interval (m_1, m_2) , we choose $m = r$. We run an MCMC chain for θ_{880} for 50000 iterations, discarding the first 1000 samples as burn-in. We

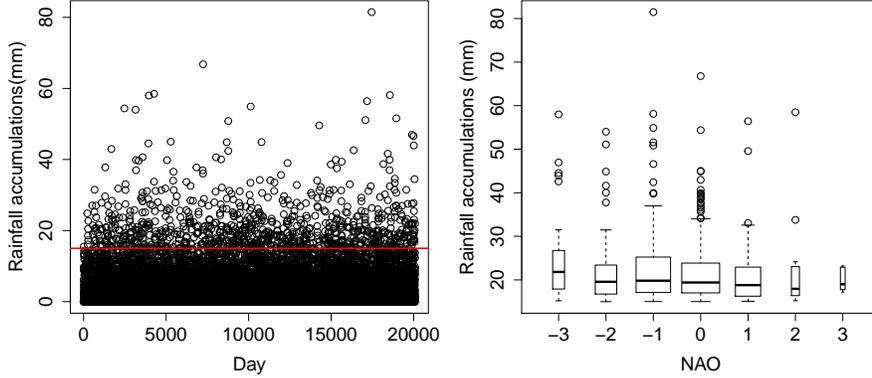


Figure 10: (Left) Daily rainfall observations in the Cumbria grid cell in the period 1958-2012. The red line represents the extreme value threshold of $u = 15$. (Right) Boxplots of rainfall against monthly NAO index.

transform the remaining samples using the mapping in (5), where $k = 55$, to obtain samples from the joint posterior of θ_{55} . The estimated posterior density for each parameter is shown in Figure 11.

To estimate probabilities of events beyond the range of the data, we can use the estimated parameters to estimate extreme quantiles of the annual maximum distribution. The quantity y_N , satisfying:

$$1/N = 1 - G(y_N), \quad (16)$$

is termed the N -year return level, where G is defined as in expression (2). The level y_N is expected to be exceeded on average once every N years. By inverting (16) we get:

$$y_N = \begin{cases} \mu_{55} - \frac{\sigma_{55}}{\xi} [1 - \{-\log(1 - 1/N)\}^{-\xi}] & \text{for } \xi \neq 0 \\ \mu_{55} - \sigma_{55} \log\{-\log(1 - 1/N)\} & \text{for } \xi = 0. \end{cases} \quad (17)$$

The posterior density of the 100-year return level in Figure 11 is estimated by inputting the MCMC samples of the model parameters into expression (17).

We use the same methodology to explore the effect of the monthly NAO index on the probability of extreme rainfall levels in Cumbria. The NAO index describes the surface sea-level pressure difference between the Azores High and the Icelandic Low. The low frequency variability of the monthly scale is chosen to represent the large scale atmospheric processes affecting the distribution of wind and rain. In the UK, a positive NAO index is associated with cool summers and wet winters, while a negative NAO index typically corresponds to cold winters, pushing the North Atlantic storm track

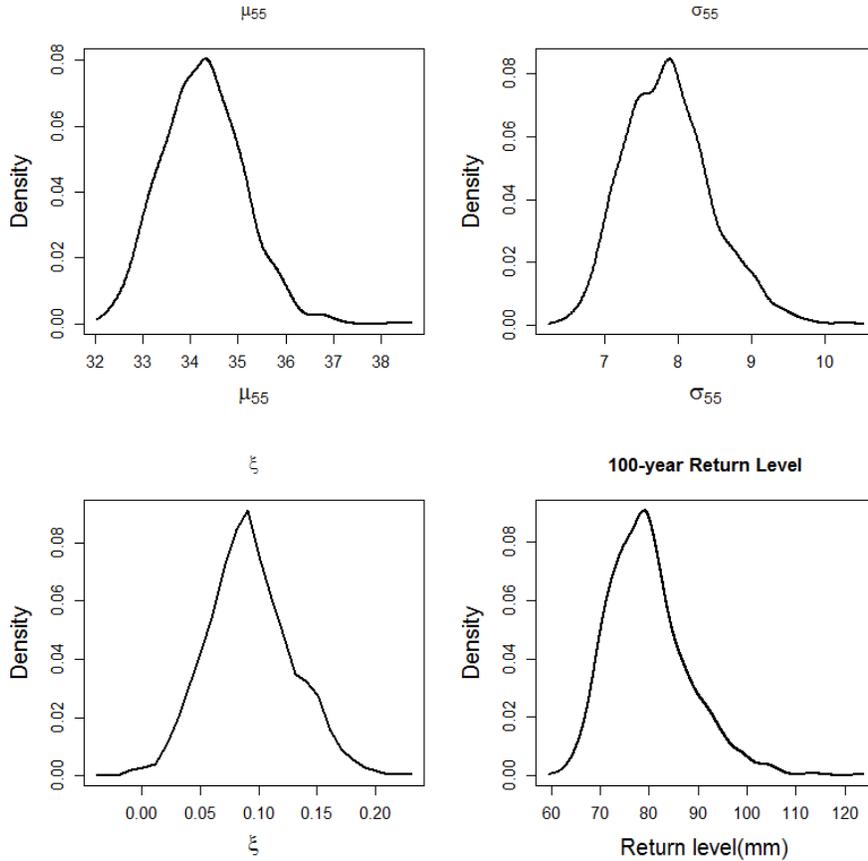


Figure 11: Estimated posterior densities of μ_{55} , σ_{55} , ξ and the 100-year return level.

further south to the Mediterranean region (Hurrell et al., 2003). In this analysis, we incorporated the effect of NAO by introducing it as a covariate in the location parameter. The threshold of $u = 15$ was retained for this analysis.

To obtain the value of m that minimises the overall correlation in the model, we solve numerically the equation $\rho_{\mu_m^{(0)}, \sigma_m} = 0$. We obtain a kernel density estimate of the NAO covariate, which represents g as defined in expression (14). We use this to obtain maximum likelihood estimates $\hat{\theta}_r(x)$. These quantities are substituted into the Fisher information matrix. The matrix is then inverted numerically to estimate $m = 920$. This represents a slight deviation from \hat{m}_2 estimated during the iid analysis. We would expect this as the covariate effect is small, as shown in Figure 12. This example illustrates the benefit of numerically solving for m when modelling non-stationarity, as the range (m_1, m_2) estimated analytically during the iid analysis no longer contain the optimal value of m .

We run an MCMC chain for $\theta_{920}(x)$ for 50000 iterations before discarding the first 1000 samples as burn-in. We transform the remaining MCMC samples to the annual maximum scale using the mapping in (15) where $k = 55$. Figure 12 indicates that NAO has a significantly positive effect on the location parameter.

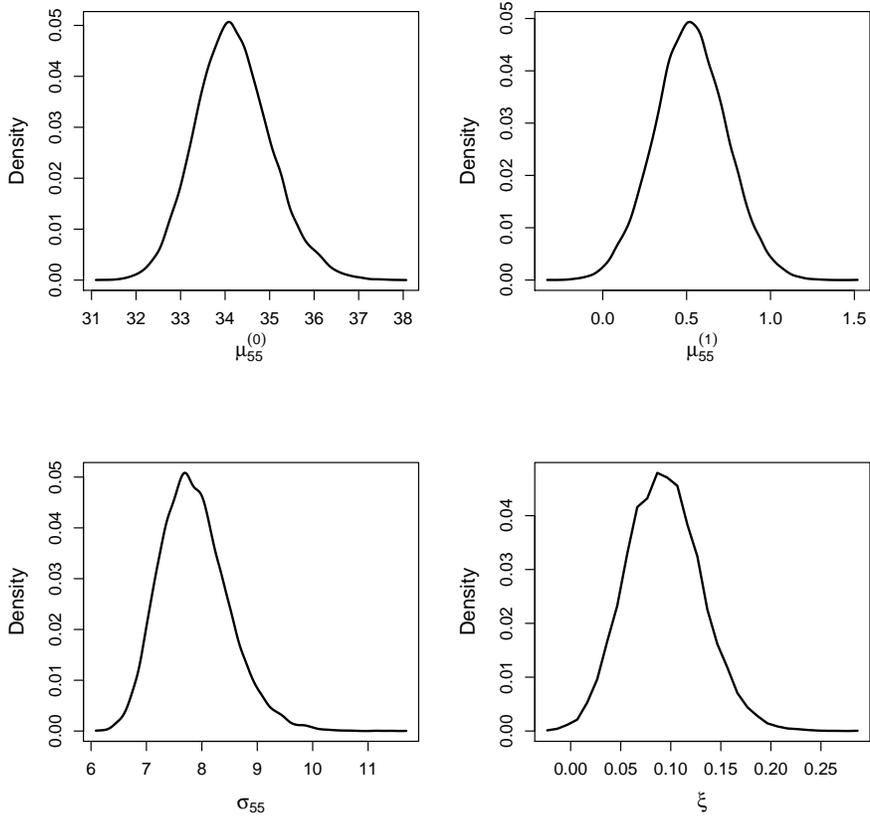


Figure 12: Estimated posterior densities of $\mu_{55}^{(0)}$, $\mu_{55}^{(1)}$, σ_{55} and ξ .

We wish to estimate return levels relating to the November 2009 flood event, which is represented by a value of 51.582mm in the dataset. Return levels corresponding to the distribution of November maxima are shown in Figure 13. We can also use the predictive distribution in order to account for both parameter uncertainty and randomness in future observations (Coles and Tawn, 1996). On the basis of threshold excesses $\mathbf{x} = (x_1, \dots, x_n)$, the predictive distribution of a future November maximum M is:

$$\Pr\{M \leq y|\mathbf{x}\} = \int_{\theta_{55}} \Pr\{M \leq y|\theta_{55}\}\pi(\theta_{55}|\mathbf{x})d\theta_{55}, \quad (18)$$

where

$$\Pr\{M \leq y | \boldsymbol{\theta}_{55}\} = \begin{cases} \exp \left\{ -\frac{1}{12} \left[1 + \xi \left(\frac{y - (\mu_{55}^{(0)} + \mu_{55}^{(1)} z)}{\sigma_{55}} \right) \right]_+^{-1/\xi} \right\} & \text{where } z \text{ is known} \\ \exp \left\{ -\frac{1}{12} \int_z \left[1 + \xi \left(\frac{y - (\mu_{55}^{(0)} + \mu_{55}^{(1)} z)}{\sigma_{55}} \right) \right]_+^{-1/\xi} g_N(z) dz \right\} & \text{where } z \text{ is unknown,} \end{cases} \quad (19)$$

where g_N is the density of NAO in November and the integral is evaluated numerically using adaptive quadrature methods. The integral in (18) can be approximated using a Monte Carlo summation over the samples from the joint posterior of $\boldsymbol{\theta}_{55}$. From this, we estimate the predictive probability of an event exceeding 51.582 in a typical November is 0.0112, with a 95% credible interval of (0.0063, 0.0185), which corresponds to an 89-year event, (54, 158). For November 2009, when an NAO index of -0.02 was measured, the probability of such an event was 0.0111, (0.0062, 0.0184), corresponding to a 90-year event, (54, 161). For the maximum observed value of NAO in November, with $\text{NAO} = 3.04$, the predictive probability of such an event is 0.0132, (0.0073, 0.0214), which corresponds to a 75-year flood event, (47, 136). This illustrates the impact that different phases of climate variability can have on the probabilities of extreme events.

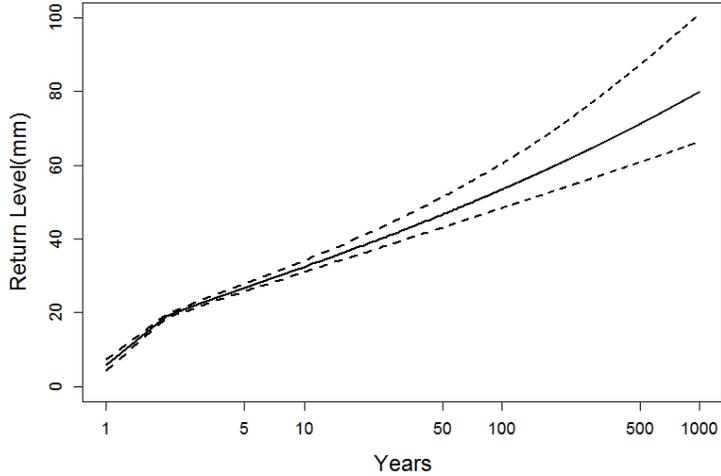


Figure 13: Return levels corresponding to November maxima. The full line represents the posterior mean and the two dashed lines representing 95% credible intervals.

Appendix

A Proof: $\hat{\mu}_r = u$ when $m = r$

We can write the full likelihood for parameters $\boldsymbol{\theta}_r$ given a series of excesses $\{x_i\}$ above a threshold u as:

$$L(\boldsymbol{\theta}_r) = L_1 \times L_2,$$

where L_1 is the Poisson probability of r exceedances of u and L_2 is the joint density of these r exceedances, so that:

$$\begin{aligned} L_1 &= \frac{1}{r!} \left\{ r \left[1 + \xi \left(\frac{u - \mu_r}{\sigma_r} \right) \right]_+^{-1/\xi} \right\}^r \exp \left\{ -r \left[1 + \xi \left(\frac{u - \mu_r}{\sigma_r} \right) \right]_+^{-1/\xi} \right\}, \\ L_2 &= \prod_{i=1}^r \frac{1}{\sigma_r} \left[1 + \xi \left(\frac{x_i - \mu_r}{\sigma_r} \right) \right]_+^{-1/\xi - 1} \left[1 + \xi \left(\frac{u - \mu_r}{\sigma_r} \right) \right]_+^{1/\xi}. \end{aligned}$$

By defining $\Lambda = \left[1 + \xi \left(\frac{u - \mu_r}{\sigma_r} \right) \right]_+^{-1/\xi}$ and using identity (A), we can reparameterise the likelihood in terms of $\boldsymbol{\theta}^* = (\Lambda, \psi_u, \xi)$ to give:

$$L(\boldsymbol{\theta}^*) \propto \Lambda^r \exp \{-r\Lambda\} \prod_{i=1}^r \frac{1}{\psi_u} \left[1 + \xi \left(\frac{x_i - u}{\psi_u} \right) \right]_+^{-1/\xi - 1}$$

Taking the log-likelihood and maximising with respect to Λ , we get:

$$\begin{aligned} l(\boldsymbol{\theta}^*) &:= \log L(\boldsymbol{\theta}^*) = r \log \hat{\Lambda} - r\hat{\Lambda} - r \log \psi_u - \left(\frac{1}{\xi} - 1 \right) \sum_{i=1}^r \log \left[1 + \xi \left(\frac{x_i - u}{\psi_u} \right) \right]_+ \\ \frac{\partial l}{\partial \Lambda} &= \frac{r}{\hat{\Lambda}} - r = 0, \end{aligned}$$

which gives $\hat{\Lambda} = 1$. Then, by the invariance property of maximum likelihood estimators, $\hat{\mu}_r = u$. Using the identity

$$\psi_u = \sigma_m + \xi(u - \mu_m),$$

which is invariant to choice of m , we get $\hat{\sigma}_r = \hat{\psi}_u$. Because the ξ -dependent term in the log-likelihood is identical to that in a GP log-likelihood, the maximum likelihood estimators of the two models coincide.

B Derivation of prior for inference on $\boldsymbol{\theta}_m$

We define a joint prior on the parameterisation of interest $\boldsymbol{\theta}_k$. However, as we are making inference for the ‘optimal’ parameterisation $\boldsymbol{\theta}_m$, we must derive the prior for $\boldsymbol{\theta}_m$. We can calculate the prior

density of $\boldsymbol{\theta}_m$ by using the density method for one-to-one bivariate transformations. Inverting (5) to get expressions for μ_m and σ_m , i.e.

$$\begin{aligned}\mu_m &= \mu_k - \frac{\sigma_k}{\xi} \left(1 - \left(\frac{m}{k} \right)^{-\xi} \right) = g_1(\mu_k, \sigma_k) \\ \sigma_m &= \sigma_k \left(\frac{m}{k} \right)^{-\xi} = g_2(\mu_k, \sigma_k),\end{aligned}$$

we can use this transformation to calculate the prior for $\boldsymbol{\theta}_m$.

$$\begin{aligned}\pi(\boldsymbol{\theta}_m) &= \pi(\mu_m, \sigma_m, \xi) \\ &= \pi(\mu_k, \sigma_k, \xi) |\det J|_{\mu_k=g_1^{-1}(\mu_m, \sigma_m), \sigma_k=g_2^{-1}(\mu_m, \sigma_m), \xi=\xi},\end{aligned}$$

where

$$\begin{aligned}\det J &= \begin{vmatrix} \frac{\partial \mu_m}{\partial \mu_k} & \frac{\partial \mu_m}{\partial \sigma_k} & \frac{\partial \mu_m}{\partial \xi} \\ \frac{\partial \sigma_m}{\partial \mu_k} & \frac{\partial \sigma_m}{\partial \sigma_k} & \frac{\partial \sigma_m}{\partial \xi} \\ \frac{\partial \xi}{\partial \mu_k} & \frac{\partial \xi}{\partial \sigma_k} & \frac{\partial \xi}{\partial \xi} \end{vmatrix} \\ &= \begin{vmatrix} \frac{\partial \mu_m}{\partial \mu_k} & \frac{\partial \mu_m}{\partial \sigma_k} & \frac{\partial \mu_m}{\partial \xi} \\ 0 & \frac{\partial \sigma_m}{\partial \sigma_k} & \frac{\partial \sigma_m}{\partial \xi} \\ 0 & 0 & \frac{\partial \xi}{\partial \xi} \end{vmatrix} \\ &= \frac{\partial \sigma_m}{\partial \sigma_k} \frac{\partial \xi}{\partial \xi} \\ &= \left(\frac{m}{k} \right)^{-\xi}.\end{aligned}$$

Therefore, $\pi(\boldsymbol{\theta}_m) = \left(\frac{m}{k} \right)^{-\xi} \pi(\boldsymbol{\theta}_k)$.

C Fisher information matrix calculations for iid random variables

The log-likelihood of the Poisson process model with parameterisation $\boldsymbol{\theta}_m = (\mu_m, \sigma_m, \xi)$ can be expressed as

$$l(\boldsymbol{\theta}_m) = -m \left[1 + \xi \left(\frac{u - \mu_m}{\sigma_m} \right) \right]_+^{-1/\xi} - r \log \sigma_m - \left(\frac{1}{\xi} + 1 \right) \sum_{j=1}^r \log \left[1 + \xi \left(\frac{x_j - \mu_m}{\sigma_m} \right) \right]_+,$$

where r is the number of exceedances of X above the threshold u . For simplicity, we drop the $[\cdot]_+$ subscript in subsequent calculations. In order to produce analytic expressions for the asymptotic covariance matrix, we must evaluate the observed information matrix $\hat{I}(\boldsymbol{\theta}_m)$. For simplicity, we define $v_m = \frac{u - \mu_m}{\sigma_m}$ and $z_{j,m} = \frac{x_j - \mu_m}{\sigma_m}$.

$$\begin{aligned}
\frac{\partial^2 l}{\partial \mu_m^2} &= -\frac{m(\xi+1)}{\sigma_m^2} [1 + \xi v_m]^{-1/\xi-2} + \frac{\xi(\xi+1)}{\sigma_m^2} \sum_{j=1}^r [1 + \xi z_{j,m}]^{-2}, \\
\frac{\partial^2 l}{\partial \sigma_m^2} &= \frac{2m}{\sigma_m^2} [1 + \xi v_m]^{-1/\xi-1} v_m - \frac{m(\xi+1)}{\sigma_m^2} [1 + \xi v_m]^{-1/\xi-2} v_m^2 + \frac{r}{\sigma_m^2} - \frac{2(\xi+1)}{\sigma_m^2} \sum_{j=1}^r [1 + \xi z_{j,m}]^{-1} z_{j,m} \\
&\quad + \frac{\xi(\xi+1)}{\sigma_m^2} \sum_{j=1}^r z_{j,m}^2 [1 + \xi z_{j,m}]^{-2}, \\
\frac{\partial^2 l}{\partial \xi^2} &= -m [1 + \xi v_m]^{-1/\xi} \left[\frac{1}{\xi} v_m^2 [1 + \xi v_m]^{-2} - \frac{2}{\xi^3} \log [1 + \xi v_m] \right. \\
&\quad \left. + \frac{2}{\xi^2} [1 + \xi v_m]^{-1} v_m + \left(\frac{1}{\xi^2} \log [1 + \xi v_m] - \frac{1}{\xi} [1 + \xi v_m]^{-1} v_m \right)^2 \right] \\
&\quad - \frac{2}{\xi^3} \sum_{j=1}^r \log [1 + \xi z_{j,m}] + \frac{2}{\xi^2} \sum_{j=1}^r [1 + \xi z_{j,m}]^{-1} z_{j,m} + \frac{\xi+1}{\xi} \sum_{j=1}^r [1 + \xi z_{j,m}]^{-2} z_{j,m}^2, \\
\frac{\partial^2 l}{\partial \mu_m \partial \sigma_m} &= \frac{m}{\sigma_m^2} [1 + \xi v_m]^{-1/\xi-1} - \frac{m(\xi+1)}{\sigma_m^2} [1 + \xi v_m]^{-1/\xi-2} v_m \\
&\quad - \frac{\xi+1}{\sigma_m^2} \sum_{j=1}^r [1 + \xi z_{j,m}]^{-1} + \frac{\xi(\xi+1)}{\sigma_m^2} \sum_{j=1}^r [1 + \xi z_{j,m}]^{-2} z_{j,m}, \\
\frac{\partial^2 l}{\partial \mu_m \partial \xi} &= -\frac{m}{\sigma_m} \left[\frac{1}{\xi^2} [1 + \xi v_m]^{-1/\xi-1} \log [1 + \xi v_m] - \frac{\xi+1}{\xi} [1 + \xi v_m]^{-1/\xi-2} v_m \right] + \frac{1}{\sigma_m} \sum_{j=1}^r [1 + \xi z_{j,m}]^{-1} \\
&\quad - \frac{\xi+1}{\sigma_m} \sum_{j=1}^r [1 + \xi z_{j,m}]^{-2} z_{j,m}, \\
\frac{\partial^2 l}{\partial \sigma_m \partial \xi} &= -\frac{m}{\sigma_m} v_m \left[\frac{1}{\xi^2} [1 + \xi v_m]^{-1/\xi-1} \log [1 + \xi v_m] - \frac{\xi+1}{\xi} [1 + \xi v_m]^{-1/\xi-2} v_m \right] + \frac{1}{\sigma_m} \sum_{j=1}^r [1 + \xi z_{j,m}]^{-1} z_{j,m} \\
&\quad - \frac{\xi+1}{\sigma_m} \sum_{j=1}^r [1 + \xi z_{j,m}]^{-2} z_{j,m}^2
\end{aligned}$$

To obtain the Fisher information matrix, we take the expected value of each term in the observed information with respect to the probability density of points of a Poisson process. Let $Z = \frac{X - \mu_m}{\sigma_m}$, and R be a random variable denoting the number of excesses of X above u . The density of points in the set A_u can be defined by

$$f(x) = \frac{\lambda(x)}{\Lambda(A_u)} = \frac{[1 + \xi z]^{-1/\xi-1}}{[1 + \xi v_m]^{-1/\xi}},$$

where λ is a function denoting the rate of exceedance. Then, for example,

$$\begin{aligned}
\mathbb{E}_{Z,R} \left\{ \sum_{j=1}^R [1 + \xi z_{j,m}]^{-2} \right\} &= \mathbb{E}_R \mathbb{E}_{Z|R} \left\{ \sum_{j=1}^R [1 + \xi z_{j,m}]^{-2} \right\} \\
&= \mathbb{E}_R \left\{ R \mathbb{E}_Z \left\{ [1 + \xi Z]^{-2} \right\} \right\} \\
&= \mathbb{E}_R \left\{ R [1 + \xi v_m]^{1/\xi} \int_{v_m}^{\infty} [1 + \xi z]^{-1/\xi-3} dz \right\} \\
&= \frac{m}{2\xi + 1} [1 + \xi v_m]^{-1/\xi-2}
\end{aligned}$$

Following this process, we can write the Fisher information matrix $I(\boldsymbol{\theta}_m)$ as:

$$\begin{aligned}
\mathbb{E} \left\{ -\frac{\partial^2 l}{\partial \mu_m^2} \right\} &= \frac{m(\xi + 1)}{\sigma_m^2} [1 + \xi v_m]^{-1/\xi-2} - \frac{m\xi(\xi + 1)}{(2\xi + 1)\sigma_m^2} [1 + \xi v_m]^{-1/\xi-2}, \\
\mathbb{E} \left\{ -\frac{\partial^2 l}{\partial \sigma_m^2} \right\} &= -\frac{2m}{\sigma_m^2} [1 + \xi v_m]^{-1/\xi-1} v_m + \frac{m(\xi + 1)}{\sigma_m^2} [1 + \xi v_m]^{-1/\xi-2} v_m^2 - \frac{r}{\sigma_m^2} + \\
&\quad \frac{2m}{\sigma_m^2} [1 + \xi v_m]^{-1/\xi-1} [1 + (\xi + 1)v_m] - \\
&\quad \frac{m\xi}{(2\xi + 1)\sigma_m^2} [1 + \xi v_m]^{-1/\xi-2} [(2\xi^2 + 3\xi + 1)v_m^2 + (4\xi + 2)v_m + 2], \\
\mathbb{E} \left\{ -\frac{\partial^2 l}{\partial \xi^2} \right\} &= m[1 + \xi v_m]^{-1/\xi} \left[\frac{1}{\xi} v_m^2 [1 + \xi v_m]^{-2} - \frac{2}{\xi^3} \log [1 + \xi v_m] + \right. \\
&\quad \left. \frac{2}{\xi^2} [1 + \xi v_m]^{-1} v_m + \left(\frac{1}{\xi^2} \log [1 + \xi v_m] - \frac{1}{\xi} [1 + \xi v_m]^{-1} \right)^2 \right] + \\
&\quad \frac{2}{\xi^3} [1 + \xi v_m]^{-1/\xi} [\xi + \log [1 + \xi v_m]] - \frac{2m}{(\xi + 1)\xi^2} [1 + \xi v_m]^{-1/\xi-1} [1 + (\xi + 1)v_m] - \\
&\quad \frac{m}{\xi(2\xi + 1)} [1 + \xi v_m]^{-1/\xi-2} [(2\xi^2 + 3\xi + 1)v_m^2 + (4\xi + 2)v_m + 2], \\
\mathbb{E} \left\{ -\frac{\partial^2 l}{\partial \mu_m \partial \sigma_m} \right\} &= \frac{m(\xi + 1)}{\sigma_m^2} [1 + \xi v_m]^{-1/\xi-2} v_m - \frac{m\xi}{(2\xi + 1)\sigma_m^2} [1 + \xi v_m]^{-1/\xi-2} [1 + (2\xi + 1)v_m], \\
\mathbb{E} \left\{ -\frac{\partial^2 l}{\partial \mu_m \partial \xi} \right\} &= \frac{m}{\sigma_m} \left[\frac{1}{\xi^2} [1 + \xi v_m]^{-1/\xi-1} \log [1 + \xi v_m] - \frac{\xi + 1}{\xi} [1 + \xi v_m]^{-1/\xi-2} v_m \right] - \\
&\quad \frac{m}{\sigma_m(\xi + 1)} [1 + \xi v_m]^{-1/\xi-1} + \frac{m}{\sigma_m(2\xi + 1)} [1 + \xi v_m]^{-1/\xi-2} [1 + (2\xi + 1)v_m], \\
\mathbb{E} \left\{ -\frac{\partial^2 l}{\partial \sigma_m \partial \xi} \right\} &= \frac{m}{\sigma_m} v_m \left[\frac{1}{\xi^2} [1 + \xi v_m]^{-1/\xi-1} \log [1 + \xi v_m] + \frac{\xi + 1}{\xi} [1 + \xi v_m]^{-1/\xi-2} v_m \right] - \\
&\quad \frac{m}{\sigma_m(\xi + 1)} [1 + \xi v_m]^{-1/\xi-1} [1 + (\xi + 1)v_m] + \\
&\quad \frac{m}{\sigma_m(2\xi + 1)} [1 + \xi v_m]^{-1/\xi-2} [(2\xi^2 + 3\xi + 1)v_m^2 + (4\xi + 2)v_m + 2].
\end{aligned}$$

By inverting the Fisher information matrix using a technical computing tool like Wolfram Mathematica, making the substitution $r = m[1 + \xi v_m]^{-1/\xi}$, the **expected** number of exceedances, and

using the mapping in (5), we can get expressions for asymptotic posterior covariances.

$$\begin{aligned} \text{ACov}(\mu_m, \xi) &= \frac{1}{\xi^2 r} (\xi + 1) \sigma_m \left(\frac{r}{m}\right)^{-\xi} \left(\xi(\xi + 1) \left(\frac{r}{m}\right)^\xi \log\left(\frac{r}{m}\right) - (2\xi + 1) \left(\left(\frac{r}{m}\right)^\xi - 1\right) \right) \\ \text{ACov}(\mu_m, \sigma_m) &= \frac{1}{\xi^2 r} \sigma_m^2 \left(\frac{r}{m}\right)^{-\xi} \left(\left(\frac{r}{m}\right)^\xi \left((\xi + 1) \log\left(\frac{r}{m}\right) \left((\xi + 1) \xi \log\left(\frac{r}{m}\right) - 3\xi - 1 \right) + \right. \right. \\ &\quad \left. \left. \xi(\xi(\xi + 2) + 3) + 1 \right) + (\xi + 1)(2\xi + 1) \left(\log\left(\frac{r}{m}\right) - 1 \right) \right) \\ \text{ACov}(\sigma_m, \xi) &= \frac{1}{r} (\xi + 1) \sigma_m \left((\xi + 1) \log\left(\frac{r}{m}\right) - 1 \right) \end{aligned}$$

When $m = r$, $\text{ACov}(\mu_m, \xi) = 0$. In addition, the m for which $\text{ACov}(\mu_m, \sigma_m) = 0$ coincides with the value of m that minimises ρ_{θ_m} as defined in (8). This root can easily be found numerically, but an analytical approximation can be calculated using a one-step Halley's method. By using $m = r$ as the initial seed, and using the formula:

$$x_{n+1} = x_n - \frac{f(x_n)}{f'(x_n) - \frac{f(x_n)f''(x_n)}{2f'(x_n)}}$$

we get the expression (13) for \hat{m}_2 after one step. The quantity for \hat{m}_1 , given by expression (12) requires two iterations of this method.

Acknowledgements

We gratefully acknowledge the support of the EPSRC funded EP/H023151/1 STOR-i Centre for Doctoral Training, the Met Office and EDF Energy. We extend our thanks to Jenny Wadsworth of Lancaster University and Simon Brown of the Met Office for helpful comments. We also thank the Met Office for the rainfall data.

References

- Attalides, N. (2015). *Threshold-based extreme value modelling*. PhD thesis, UCL (University College London).
- Chavez-Demoulin, V. and Davison, A. C. (2005). Generalized additive modelling of sample extremes. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 54(1):207–222.
- Coles, S. G. (2001). *An Introduction to Statistical Modeling of Extreme Values*. Springer.
- Coles, S. G. and Tawn, J. A. (1996). A Bayesian analysis of extreme rainfall data. *Applied Statistics*, 45:463–478.

- Cox, D. R. and Reid, N. (1987). Parameter orthogonality and approximate conditional inference (with discussion). *Journal of the Royal Statistical Society. Series B (Methodological)*, 49(1):1–39.
- Davison, A. C. and Smith, R. L. (1990). Models for exceedances over high thresholds (with discussion). *Journal of the Royal Statistical Society. Series B (Methodological)*, 52(3):393–442.
- Efron, B. and Hinkley, D. V. (1978). Assessing the accuracy of the maximum likelihood estimator: Observed versus expected fisher information. *Biometrika*, 65(3):457–483.
- Gander, W. (1985). On Halley’s iteration method. *American Mathematical Monthly*, 92(2):131–134.
- Hills, S. E. and Smith, A. F. (1992). Parameterization issues in Bayesian inference. *Bayesian Statistics*, 4:227–246.
- Hurrell, J. W., Kushnir, Y., Ottersen, G., and Visbeck, M. (2003). An overview of the North Atlantic oscillation. *Geophysical Monograph-American Geophysical Union*, 134:1–36.
- Northrop, P. J. and Attalides, N. (2016). Posterior propriety in Bayesian extreme value analyses using reference priors. *Statistica Sinica*, 26(2):721–743.
- Northrop, P. J. and Jonathan, P. (2011). Threshold modelling of spatially dependent non-stationary extremes with application to hurricane-induced wave heights. *Environmetrics*, 22(7):799–809.
- Pickands, J. (1975). Statistical inference using extreme order statistics. *The Annals of Statistics*, 3(1):119–131.
- Robert, C. and Casella, G. (2009). *Introducing Monte Carlo Methods with R*. Springer Science & Business Media.
- Roberts, G. O., Rosenthal, J. S., et al. (2001). Optimal scaling for various Metropolis-Hastings algorithms. *Statistical Science*, 16(4):351–367.
- Smith, R. L. (1985). Maximum likelihood estimation in a class of nonregular cases. *Biometrika*, 72(1):67–90.
- Smith, R. L. (1987a). Discussion of “Parameter orthogonality and approximate conditional inference” by D.R. Cox and N. Reid. *Journal of the Royal Statistical Society. Series B (Methodological)*, 49(1):21–22.
- Smith, R. L. (1987b). A theoretical comparison of the annual maximum and threshold approaches to extreme value analysis. Technical Report 53, University of Surrey.

- Smith, R. L. (1989). Extreme value analysis of environmental time series: an application to trend detection in ground-level ozone. *Statistical Science*, 4(4):367–377.
- Stephenson, A. (2016). Bayesian inference for extreme value modelling. *Extreme Value Modeling and Risk Analysis: Methods and Applications*, pages 257–280.
- Tawn, J. A. (1987). Discussion of “Parameter orthogonality and approximate conditional inference” by D.R. Cox and N. Reid. *Journal of the Royal Statistical Society. Series B (Methodological)*, 49(1):33–34.
- Wadsworth, J. L., Tawn, J. A., and Jonathan, P. (2010). Accounting for choice of measurement scale in extreme value modeling. *The Annals of Applied Statistics*, 4(3):1558–1578.