



HAL
open science

Reparameterization of extreme value framework for improved Bayesian workflow

Théo Moins, Julyan Arbel, Stéphane Girard, Anne Dutfoy

► **To cite this version:**

Théo Moins, Julyan Arbel, Stéphane Girard, Anne Dutfoy. Reparameterization of extreme value framework for improved Bayesian workflow. *Computational Statistics and Data Analysis*, 2023, 187, pp.1-21. 10.1016/j.csda.2023.107807 . hal-03806159v4

HAL Id: hal-03806159

<https://hal.science/hal-03806159v4>

Submitted on 9 Jun 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Reparameterization of extreme value framework for improved Bayesian workflow

Théo Moins * Julyan Arbel * Stéphane Girard * Anne Dutfoy †

June 9, 2023

Abstract

Using Bayesian methods for extreme value analysis offers an alternative to frequentist ones, with several advantages such as easily dealing with parametric uncertainty or studying irregular models. However, computations can be challenging and the efficiency of algorithms can be altered by poor parametrization choices. The focus is on the Poisson process characterization of univariate extremes and outline two key benefits of an orthogonal parameterization. First, Markov chain Monte Carlo convergence is improved when applied on orthogonal parameters. This analysis relies on convergence diagnostics computed on several simulations. Second, orthogonalization also helps deriving Jeffreys and penalized complexity priors, and establishing posterior propriety thereof. The proposed framework is applied to return level estimation of Garonne flow data (France).

1 Introduction

Studying the long-term behavior of environmental variables is necessary to understand the risks of hazardous meteorological events such as floods, storms, or droughts. To this end, models from extreme value theory allow us to extrapolate data in the distribution tails, in order to estimate extreme quantiles that may not have been observed (see [Coles, 2001](#), for an introduction). In particular, a key quantity to estimate is the return level ℓ_T associated with a given period of T years, the level that is exceeded on average once every T years. Assessing the resistance of facilities to natural disasters such as dams to floods that occur on average once every 100 years or 1000 years is critical for companies such as *Électricité de France* (EDF). Moreover, characterizing the uncertainty on the estimation of this return level is also of interest, which encourages the choice of the Bayesian paradigm. However, performing Bayesian inference requires multiple steps that must be managed by the user, from the choice of the model to the evaluation and validation of computations. This has been recently formalized by [Gelman et al. \(2020\)](#) in the form of a Bayesian workflow. After introducing models stemming from extreme value theory in [Section 1.1](#), we briefly review in [Section 1.2](#) one particular step of the workflow, reparameterization, and more specifically the choice of an orthogonal parameterization.

*Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000 Grenoble, France.

†EDF R&D dept. Périclès, 91120 Palaiseau, France.

1.1 Extreme-value models

Three different frameworks exist to model extreme events, leading to different likelihoods: one by block maxima, one by peaks-over-threshold, and one that unifies both through a Poisson process characterization.

Block maxima model Let M_n be the maximum of n i.i.d random variables with cumulative distribution function (cdf) F . We assume that F belongs to the maximum domain of attraction of a non-degenerate cdf G , meaning that there exist two sequences $a_n > 0$ and b_n such that $(M_n - b_n)/a_n$ converges in distribution to the cdf G . The extreme value theorem (e.g., [Haan and Ferreira, 2006](#), Chapter 1) states that G is necessarily a generalized extreme-value (GEV) distribution, with cdf:

$$G(x) = \begin{cases} \exp\left(-\{1 + \xi x\}_+^{-1/\xi}\right) & \text{if } \xi \neq 0, \\ \exp(-\exp(-x)) & \text{if } \xi = 0, \end{cases} \quad (1)$$

where $\{x\}_+ = \max\{0, x\}$. Consequently, for a finite value of n , one can consider the approximation $\mathbb{P}(M_n \leq x) \approx G((x - b_n)/a_n) =: G(x | b_n, a_n, \xi)$, and focus on the estimation of the three parameters of the GEV distribution. Here, as the dataset is fixed, the dependence in n for the location and scale parameters will be omitted. To obtain a sample of maxima, one can divide the dataset into m blocks of size n/m and extract the maximum from each of them.

Peaks-over-threshold model Alternatively, one can consider observations that exceed a high threshold u . Let X be a random variable with cdf F . Pickands theorem ([Pickands, 1975](#)) states that, if F belongs to the maximum domain of attraction of G with $\mathbb{P}(M_n \leq x) \approx G(x | \mu, \sigma, \xi)$, then the distribution of the exceedances $X - u | X > u$ is, as u converges to the upper endpoint of F , a generalized Pareto distribution (GPD), with cdf

$$H(y | \tilde{\sigma}, \xi) = \begin{cases} 1 - \{1 + \xi \frac{y}{\tilde{\sigma}}\}_+^{-1/\xi} & \text{if } \xi \neq 0, \\ 1 - \exp\left(-\frac{y}{\tilde{\sigma}}\right) & \text{if } \xi = 0, \end{cases} \quad (2)$$

where the shape parameter ξ is the same as in (1) and the GPD and GEV scales are linked by $\tilde{\sigma} = \sigma + \xi(u - \mu)$. To obtain a sample of n_u excesses, the peaks-over-threshold method focusses on the n_u largest values of the dataset. It thus requires the estimation of the quantile of order $1 - n_u/n$, which can be seen as the third parameter to estimate, in addition to $\tilde{\sigma}$ and ξ . The most classical choice is to estimate this intermediate quantile by the $(n - n_u)$ th order statistic.

Poisson process characterization of extremes Finally, these two approaches can be generalized by a third one, using a non-homogeneous Poisson process. We present here an intuitive way for obtaining this model similarly to [Coles \(2001, Chapter 7\)](#), and refer to ([Leadbetter et al., 1983](#), Chapter 5) for theoretical details. We start by observing that, for large n , $F^n(x) \approx G(x | \mu, \sigma, \xi)$, for x in the support of G denoted by $\text{supp}(G(\cdot | \mu, \sigma, \xi)) = \{x \in \mathbb{R} \text{ s.t. } 1 + \xi \left(\frac{x - \mu}{\sigma}\right) > 0\}$. Hence, considering a large threshold $u \in \text{supp}(G(\cdot | \mu, \sigma, \xi))$, a Taylor expansion yields

$$n \log F(u) \simeq -n(1 - F(u)) \simeq \log G(u | \mu, \sigma, \xi),$$

or, equivalently,

$$\mathbb{P}(X > u) \simeq -\frac{1}{n} \log G(u | \mu, \sigma, \xi). \quad (3)$$

Equation (3) can be seen as the probability of X to belong to $I_u := [u, +\infty)$. In the case of n i.i.d random variables, one can deduce that the associated point process N_n is such that $N_n(I_u) \sim \mathcal{B}(n, p_n)$ with p_n given by Equation (3). As $n \rightarrow \infty$, the binomial distribution $\mathcal{B}(n, p_n)$ converges to the Poisson distribution $\mathcal{P}(\Lambda(I_u))$, with $\Lambda(I_u) = -\log G(u | \mu, \sigma, \xi)$. This property being valid for all I_u together with the independence property on non-overlapping sets imply that N_n converges to a non-homogeneous Poisson process, with intensity measure $\Lambda(I_u)$: $N_n \xrightarrow{d} N$, with $N(I_u) \sim \mathcal{P}(\Lambda(I_u))$. This model generalizes the block maxima one since

$$\mathbb{P}(M_n < x) = \mathbb{P}(N_n(I_x) = 0) \rightarrow \mathbb{P}(N(I_x) = 0) = \exp(-\Lambda(I_x)) = G(x | \mu, \sigma, \xi),$$

as $n \rightarrow \infty$. However, an estimation of the parameters (μ, σ, ξ) with this model is related to the overall maximum M_n of the dataset, and it is frequent to study maxima of m smaller blocks $M_{n/m}$, where m is typically the number of years in the observations and so $M_{n/m}$ corresponds to annual maxima. To do so, the intensity measure is multiplied by m , which modifies the parameterization and in particular the value of μ and σ : [Wadsworth et al. \(2010\)](#) shows that, if $(\mu_{k_i}, \sigma_{k_i}, \xi)$ ($i \in \{1, 2\}$), are parameters for k_i GEV observations, then

$$\mu_{k_2} = \mu_{k_1} - \frac{\sigma_{k_1}}{\xi} \left(1 - \left(\frac{k_2}{k_1} \right)^{-\xi} \right), \quad \sigma_{k_2} = \sigma_{k_1} \left(\frac{k_2}{k_1} \right)^{-\xi}. \quad (4)$$

The threshold excess model can also be derived from the point process representation, since $\mathbb{P}(X > y + u | X > u) \simeq 1 - H(y | \tilde{\sigma}, \xi)$, with $\tilde{\sigma} = \sigma + \xi(u - \mu)$. Moreover, in contrast to the peaks-over-threshold model where an intermediate quantile needs to be estimated, the Poisson model directly includes a third location parameter μ .

In the following, we will focus mainly on this latter model, and treat the peaks-over-threshold method as a special case in [Section 4.1](#).

Bayesian inference Using the Bayesian paradigm in extreme value models is advantageous in comparison to the frequentist approach, see [Coles and Powell \(1996\)](#) for a general review, and [Stephenson \(2016\)](#) or [Bousquet \(2021\)](#) for more recent overviews. For the Poisson process characterization of extremes, Bayesian inference consists in fixing a scaling factor m and a threshold u to get a number of $n_u \geq 1$ observations exceeding u denoted by $\mathbf{x} = (x_1, \dots, x_{n_u})$. The likelihood of these observations can be written as

$$L(\mathbf{x}, n_u | \mu, \sigma, \xi) = e^{-m(1+\xi(\frac{u-\mu}{\sigma}))^{-1/\xi}} \sigma^{-n_u} \prod_{i=1}^{n_u} \left(1 + \xi \left(\frac{x_i - \mu}{\sigma} \right) \right)^{-1-1/\xi}. \quad (5)$$

A complete Bayesian model requires also the specification of a prior $p(\mu, \sigma, \xi)$, to obtain the posterior $p(\mu, \sigma, \xi | \mathbf{x}, n_u)$ using Bayes' theorem, $p(\mu, \sigma, \xi | \mathbf{x}, n_u) \propto p(\mu, \sigma, \xi) L(\mathbf{x}, n_u | \mu, \sigma, \xi)$. This posterior summarizes the information on the parameters after observations, and can be used to extract point estimators, build credible intervals, or write the probability of a new observation \tilde{x} given data \mathbf{x} using the posterior predictive:

$$p(\tilde{x} | \mathbf{x}, n_u) = \int p(\tilde{x} | \boldsymbol{\theta}) p(\boldsymbol{\theta} | \mathbf{x}, n_u) d\boldsymbol{\theta}, \quad \boldsymbol{\theta} = (\mu, \sigma, \xi). \quad (6)$$

These quantities of interest are rarely explicit, and are often derived by sampling approaches. A recent survey of extreme value softwares ([Belzile et al., 2022](#)) contains a Bayesian section, and a comparison

with frequentist methods. In the general Bayesian case, an overview of the Bayesian workflow is given in [Gelman et al. \(2020\)](#), and we focus here on the particular step of reparameterization for the likelihood $L(\boldsymbol{x}, n_u \mid \mu, \sigma, \xi)$ in the case where Markov chain Monte Carlo (MCMC) methods are used to approximate the posterior distribution.

1.2 Reparameterization

Although the choice of parameterization of a statistical model does not alter the model *per se*, it does reshape its geometry, which in turn may impact computational aspects of sampling algorithms such as efficiency or accuracy. For these methods, a crucial complication for chain convergence is parameter correlation. This notion of correlation between parameters can be associated with a notion of asymptotic orthogonality, leading to independence of posterior components.

Parameterization and Bayesian inference It has been known for several decades that parameterization is crucial for good mixing of MCMC chains, especially when the correlation between the coordinates is large. See [Gilks et al. \(1995, Chapter 6\)](#) for a great introduction for Gibbs sampling and Metropolis–Hastings algorithm. More general computations are conducted by [Roberts and Sahu \(1997\)](#) in the normal case, but this convergence rate is less explicit in the general case, see for example [Roberts and Polson \(1994\)](#). For Metropolis–Hastings, if the structure of the kernel is not similar to the one of the target density (which is a typical case if there is a complex dependence between parameters), then too many candidates generated by the kernel are rejected and the same problem as for Gibbs sampling occurs. For more recent MCMC algorithms such as Hamiltonian Monte Carlo (HMC, [Neal, 2011](#)) and its variant NUTS ([Hoffman and Gelman, 2014](#)), [Betancourt and Girolami \(2015\)](#) gives an example of the benefit of reparameterization for hierarchical models. More generally, [Betancourt \(2019\)](#) studies reparameterization from a geometric perspective, in order to show its equivalence with adapted versions of HMC on Riemannian manifolds.

Due to the difficulty of obtaining general results on reparameterization and MCMC convergence, a significant part of the research focuses on specific models, such as hierarchical models ([Papaspiliopoulos et al., 2003](#); [Browne et al., 2009](#)), linear regression ([Gilks et al., 1995](#)), or mixed models ([Gelfand et al., 1995, 1996](#)).

For extreme value models, [Diebolt et al. \(2005\)](#) uses a continuous mixture of exponential distributions in the GPD case. [Opitz et al. \(2018\)](#) also suggests to use the median instead of the usual scale parameter to reduce correlation for Integrated Nested Laplace approximation (INLA). An alternative Monte Carlo algorithm, the ratio-of-uniforms method, is also implemented for extreme value models in the `revdbayes` package ([Northrop, 2023](#)). The influence of parameterization is also considered in this framework as the acceptance rate can be altered because of correlated parameters (see [Appendix C.4](#)). Parameter transformations are also studied in order to make likelihood-based inference suitable in the high-dimensional case in [Jóhannesson et al. \(2022\)](#). Finally, [Belzile et al. \(2022\)](#) proposes a reparameterization trick that can be used to obtain a suitable initial value for optimization routines.

Orthogonal parameterization As seen before, reducing dependence between coordinates is desirable for MCMC methods. Dependence can be characterized using asymptotic covariance and the notion of orthogonality according to [Jeffreys \(1961\)](#): parameters are said to be orthogonal when the Fisher information is diagonal. From this definition, having orthogonal parameters leads to asymptotic posterior independence when a Bernstein–von Mises theorem holds (e.g., [Van der Vaart,](#)

2000, Chapter 10). However, the problem of finding an orthogonal parameterization is seldom feasible when there are more than three parameters, since the number of equations is then greater than the number of unknown variables. In the case of three parameters, there are as many equations as there are unknowns, but the non linear system does not necessarily lead to a solution (Huzurbazar, 1950).

The main use of orthogonal parameterization is to make parameters of interest independent of nuisance parameters (Cox and Reid, 1987). Other definitions of orthogonality are also proposed to be more adapted to the inferential context (Tibshirani and Wasserman, 1994) or to ensure consistency of the parameter of interest (Woutersen, 2011). For Bayesian inference, Tibshirani and Wasserman (1994) compares different definitions and suggests a strong assumption of normality for the posterior. In the following, we keep the most popular definition of orthogonality due to Jeffreys (1961), as we are not interested in properties associated with the estimation of a given parameter of interest, but rather on the dependence structure between parameters. However, up to our knowledge, there is no clear evidence in the literature of a direct link between parameter orthogonality and mixing properties of the corresponding MCMC chains, such as a better convergence rate. In Section 4, we bring some empirical evidence on the interest of orthogonality in extreme value models.

1.3 Contributions and outline

In this paper, we study the benefits of reparameterization for the Poisson process characterization of extremes in a Bayesian context. In particular, it is shown that the orthogonal parameterization is useful for several reasons: we argue in Section 2 that it improves the performance of MCMC algorithms in terms of convergence, and we show in Section 3 that it also facilitates the derivation of priors such as Jeffreys and an informative variant on the shape parameter using penalized complexity (PC) priors (Simpson et al., 2017). These results are then illustrated by experiments in Section 4, first on simulations to compare the different parameterizations, and second on a real dataset of the Garonne river flow. Proofs as well as additional experiments are provided in the Appendix, and the code corresponding to the experiments is available online.¹

2 Reaching orthogonality for extreme Poisson process

An attempt to reparametrize the Poisson process for extremes in order to improve MCMC convergence already exists in the literature (Sharkey and Tawn, 2017), but has several limitations that are detailed here. Instead, we suggest to use the fully orthogonal parameterization of Chavez-Demoulin and Davison (2005).

Near-orthogonality with hyperparameter tuning Based on the relationship between parameters given in Equation (4), Sharkey and Tawn (2017) suggests to change the scaling factor m before using Metropolis–Hastings algorithm in order to optimize MCMC convergence. To this aim, they minimize the non-diagonal elements of the inverse Fisher information matrix corresponding to asymptotic covariances and then retrieved the parameters corresponding to the initial number of blocks from Equation (4). As the calculations cannot be achieved explicitly, the authors found empirically that the values m_1 and m_2 that cancel respectively the asymptotic covariances $\text{ACov}(\mu, \sigma)$ and $\text{ACov}(\sigma, \xi)$ are such that any $m \in [m_1, m_2]$ improves the MCMC convergence. Approximations

¹<https://github.com/TheoMoins/ExtremesPyMC>

of m_1 and m_2 are then given as functions of ξ , and therefore a preliminary estimation of ξ (typically obtained using maximum likelihood estimation) is required to obtain $\hat{m}_1(\xi)$ and $\hat{m}_2(\xi)$, and to choose a value in this interval before running an MCMC with the right choice of m . Despite leading to significant improvement of the convergence of Markov Chains, this method suffers from several limitations. First, preliminary estimation of the shape parameter ξ is required, to compute $\hat{m}_1(\xi)$ and $\hat{m}_2(\xi)$ and choose a value in the corresponding interval, which adds complexity and computational burden on the overall framework. Moreover, it also affects the accuracy of orthogonalization, as the expressions of m_1 and m_2 are found empirically, then are approximated by $\hat{m}_1(\xi)$ and $\hat{m}_2(\xi)$, and finally computed at $\hat{\xi}$ which adds a new source of uncertainty. One way to lighten the method would be to suggest a simpler choice of m , for example $m = n_u$, which leads to a satisfactory behaviour as noticed by [Wadsworth et al. \(2010\)](#). However, it is shown in [Appendix A](#) that this choice presents some flaws and does not bring any general guarantee of orthogonality.

Orthogonal parameterization Alternatively, there exists a parameterization of the Poisson process that leads to orthogonality. Suggested by [Chavez-Demoulin and Davison \(2005\)](#), it consists of the change of variable

$$(r, \nu, \xi) = \left(m \left(1 + \xi \left(\frac{u - \mu}{\sigma} \right) \right)^{-1/\xi}, (1 + \xi)(\sigma + \xi(u - \mu)), \xi \right), \quad (7)$$

while the inverse transformation is

$$(\mu, \sigma, \xi) = \left(u - \frac{\nu}{\xi(1 + \xi)} \left(1 - \left(\frac{r}{m} \right)^\xi \right), \frac{\nu}{(1 + \xi)} \left(\frac{r}{m} \right)^\xi, \xi \right).$$

With this parameterization, the likelihood is

$$L(\mathbf{x}, n_u | r, \nu, \xi) = e^{-r} \left(\frac{r}{m} \right)^{n_u} \left(\frac{\nu}{1 + \xi} \right)^{-n_u} \prod_{i=1}^{n_u} \left(1 + \frac{\xi(1 + \xi)}{\nu} (x_i - u) \right)^{-1-1/\xi}. \quad (8)$$

Under this form, we can directly see that r is orthogonal to ν and ξ , as the likelihood factorizes with respect to r and (ν, ξ) . Parameter $r \geq 0$ represents the intensity of the Poisson process, which is the expected number of exceedances, while the two other ones can be seen as an orthogonal parameterization of the GPD with scale $\tilde{\sigma}_u = \sigma + \xi(u - \mu)$ and shape ξ . Under this parameterization and provided $\xi > -1/2$, the Fisher information matrix $\mathcal{I}(r, \nu, \xi)$ is

$$\mathcal{I}(r, \nu, \xi) = \text{diag} \left(\frac{1}{r}, \frac{r}{\nu^2(1 + 2\xi)}, \frac{r}{(1 + \xi)^2} \right), \quad (9)$$

where $\text{diag}(\mathbf{u})$ denotes the diagonal matrix with diagonal equal to vector \mathbf{u} . Calculations are provided in [Appendix B](#). Therefore, the orthogonal parameterization of [Chavez-Demoulin and Davison \(2005\)](#) is more adapted than the tuning of m since it directly yields the optimal solution sought by [Sharkey and Tawn \(2017\)](#). Moreover, it is obtained without recourse to any optimization procedure or approximation. Finally, by plugging (r, ν) into Equation (4), one can show that the invariance property with respect to m holds for the three parameters, and so the parameterization is independent of the choice of m .

Generalisation to covariates The inclusion of covariates in a model holds both theoretical and practical significance, as it enables the incorporation of factors such as temporal trends. A notable advantage of this approach is that if the parameters are orthogonal and each of them depends on distinct parameters, then these parameters will also be orthogonal to one another. To elaborate further, let \mathbf{C}_i represent the set of covariates associated with the observation x_i . In the most comprehensive scenario, where all relevant covariates are considered, the model can be expressed as follows:

$$\begin{cases} r_i = f_r(\boldsymbol{\theta}_r, \mathbf{C}_i), \\ \nu_i = f_\nu(\boldsymbol{\theta}_\nu, \mathbf{C}_i), \\ \xi_i = f_\xi(\boldsymbol{\theta}_\xi, \mathbf{C}_i), \end{cases}$$

where $\boldsymbol{\theta}_r$, $\boldsymbol{\theta}_\nu$ and $\boldsymbol{\theta}_\xi$ are three vectors of parameters. The log-likelihood can be written as

$$\sum_{i=1}^n \ell_i(r_i, \nu_i, \xi_i) = \sum_{i=1}^n \ell_i(f_r(\boldsymbol{\theta}_r, \mathbf{C}_i), f_\nu(\boldsymbol{\theta}_\nu, \mathbf{C}_i), f_\xi(\boldsymbol{\theta}_\xi, \mathbf{C}_i)).$$

By leveraging the property that no parameters are shared, along with the chain rule, one can derive the following expression for the Fisher information associated with the j th coordinate $\theta_{\nu,j}$ of $\boldsymbol{\theta}_\nu$ and the k th coordinate $\theta_{\xi,k}$ of $\boldsymbol{\theta}_\xi$:

$$\mathbb{E} \left(-\frac{\partial^2 \ell_i}{\partial \theta_{\nu,j} \partial \theta_{\xi,k}} \right) = \frac{\partial f_\nu}{\partial \theta_{\nu,j}} \frac{\partial f_\xi}{\partial \theta_{\xi,k}} \mathbb{E} \left(-\frac{\partial^2 \ell_i}{\partial \nu_i \partial \xi_i} \right) = 0.$$

As a result, every parameter in $\boldsymbol{\theta}_\nu$ is orthogonal to all parameters in $\boldsymbol{\theta}_\xi$. Similar calculations reveal that $\boldsymbol{\theta}_r$ is orthogonal to both $\boldsymbol{\theta}_\nu$ and $\boldsymbol{\theta}_\xi$. Hence, when covariates are defined based on orthogonal parameters, it leads to block-wise orthogonality, with the associated Fisher information matrix comprising three blocks in this case.

3 Priors invariant to reparameterization

In the case where no external information is available about the parameters, the choice of the prior distribution should be made with caution. Typically, the term “uninformative prior” or “objective prior” can be misleading, as it refers to priors used when one does not have preliminary information, but the prior itself does contain information. As an example, a flat prior over the range of possible values is only flat for a given parameterization. After a change of variable, a uniform prior does not necessarily remain uniform (e.g., [Robert, 2007](#), Chapter 3). This problem is all the more serious as our study which deals with reparameterization: here, we derive two priors that enjoy the property of being invariant with respect to reparameterization.

3.1 Jeffreys prior

Jeffreys prior ([Jeffreys, 1946](#)) is built with the aim of invariance: if $\mathcal{I}(\boldsymbol{\theta})$ denotes the Fisher information matrix associated with parameters $\boldsymbol{\theta}$, it is defined as

$$p_J(\boldsymbol{\theta}) \propto \sqrt{\det \mathcal{I}(\boldsymbol{\theta})}. \quad (10)$$

Under this prior, one can show that a reparameterization $\boldsymbol{\phi} = h(\boldsymbol{\theta})$ with h a continuously differentiable function yields $p_J(\boldsymbol{\phi}) \propto \sqrt{\det \mathcal{I}(\boldsymbol{\phi})}$. This prior is computed for the GPD by [Castellanos and Cabras](#)

(2007) and for the GEV under a modified version where $p_J(\mu, \sigma, \xi) \propto \sqrt{\det \mathcal{I}(\sigma, \xi)}$ by [Kotz and Nadarajah \(2000\)](#). Up to our knowledge, Jeffreys prior has never been computed for the Poisson process characterization of extremes. The orthogonalization done in Equation (7) directly provides Jeffreys prior with respect to (r, ν, ξ) :

Proposition 1 *Jeffreys prior associated with a Poisson process for extremes with parameters (r, ν, ξ) from Equation (7) exists provided $\xi > -1/2$, and is given by*

$$p_J(r, \nu, \xi) \propto \frac{r^{1/2}}{\nu(1+\xi)(1+2\xi)^{1/2}}. \quad (11)$$

Moreover, the invariance to reparameterization property directly provides the expression of Jeffreys prior on (μ, σ, ξ) .

Corollary 1 *Jeffreys prior associated with a Poisson process for extremes with original parameters (μ, σ, ξ) exists provided $\xi > -1/2$, and can be written as*

$$p_J(\mu, \sigma, \xi) \propto \frac{\left(1 + \xi \left(\frac{u-\mu}{\sigma}\right)\right)^{-\frac{3}{2\xi}-1}}{\sigma^2(1+\xi)(1+2\xi)^{1/2}}. \quad (12)$$

This prior cannot be defined for $\xi \leq -1/2$, as it corresponds to a case where the Fisher information matrix is infinite. However, this assumption is not too restrictive as the great majority of models of interest belong to a maximum domain of attraction with $\xi \in (-1/2, 1/2)$. Note that this prior, similarly to the uniform one, is improper in the sense that the integral over the range of parameters is infinite. Consequently, it is necessary to check whether the posterior is proper or not to be able to use it. [Castellanos and Cabras \(2007\)](#) shows that the posterior is proper when using Jeffreys prior in the GPD case, while [Northrop and Attalides \(2016\)](#) shows that it is never the case with GEV likelihood. For the Poisson process, we show the following result:

Proposition 2 *Jeffreys prior for a Poisson process for extremes yields a proper posterior distribution, as long as $\xi > -1/2$.*

A proof is provided in [Appendix B](#).

3.2 Penalized complexity prior for the shape parameter

The shape parameter ξ plays a crucial role in the inference, as it tunes the heaviness of the tail distribution: it is heavy if $\xi > 0$, light if $\xi = 0$ and finite (*i.e.* with a finite right end-point) if $\xi < 0$. The case $\xi = 0$ can be seen as a simpler model with an exponential decrease of the survival function, where the GPD cdf in Equation (2) simplifies to an exponential distribution. This concentration of an entire maximum domain of attraction at a single value of ξ complicates the study, since it is for example difficult to distinguish heavy tails with low ξ from light tails ([Stephenson and Tawn, 2004](#)). However, this change of regime can have significant consequences when it comes to extrapolation. It should also be noted that a vast majority of datasets have distribution with $|\xi| \leq 1/2$. It is therefore natural, even in a non-informative framework, to penalize high values of $|\xi|$. One way to do this is to use penalized complexity (PC) priors ([Simpson et al., 2017](#)): the idea is to consider a prior that penalizes exponentially the distance between a model $p_\xi := p(\cdot | \xi)$ with a given ξ and the baseline p_0 with $\xi = 0$. The general formula is

$$p_{\text{PC}}(\xi | \lambda) = \lambda \exp(-\lambda d(\xi)) \left| \frac{\partial d(\xi)}{\partial \xi} \right|,$$

with $\lambda > 0$, $d(\xi) = \sqrt{2\text{KL}(p_\xi||p_0)}$ and $\text{KL}(p_\xi||p_0)$ the Kullback–Leibler divergence between p_ξ and p_0 : $\text{KL}(p_\xi||p_0) = \int p_\xi(x) \log(p_\xi(x)/p_0(x)) dx$. Parameter λ acts as a scaling parameter and controls the range of acceptable values for ξ . This prior has the advantage of being proper and invariant to reparameterization on ξ . The computation with GPD has already been done by [Opitz et al. \(2018\)](#) for the case $\xi \geq 0$: the authors prove that $d(\xi)$ is finite only if $\xi < 1$, and is $d(\xi) = \sqrt{2\xi/\sqrt{1-\xi}}$ for $0 \leq \xi < 1$. Then, they show that it can be approximated by an exponential distribution on ξ in the case $\xi \rightarrow 0$, when λ can be taken sufficiently large and to favor $\xi = 0$. A first observation is that routine calculations extend this definition both to negative values of ξ , and to the Poisson process characterization where the density of observation is also GPD.

Proposition 3 *The PC prior associated with a Poisson process for extremes exists for any $\xi < 1$ and is*

$$p_{\text{PC}}(\xi | \lambda) = \frac{\lambda}{2} \left(\frac{1 - \xi/2}{(1 - \xi)^{3/2}} \right) \exp \left(-\lambda \frac{|\xi|}{\sqrt{1 - \xi}} \right). \quad (13)$$

This prior is plotted for several values of λ in [Figure 1](#). As observed by [Opitz et al. \(2018\)](#), the PC prior is very similar to a $\text{Laplace}(0, 1/\lambda)$ when λ is large enough so that the peaks at 0 dominates over the endpoint at 1. In the case where λ is small (typically $\lambda \leq 1$), an asymptote appears at the

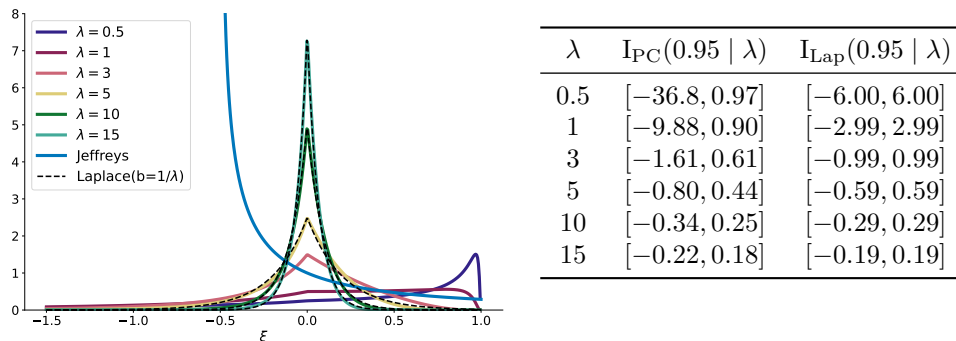


Figure 1: Left panel: examples of PC priors $p_{\text{PC}}(\cdot | \lambda)$ with λ ranging from 0.5 to 15, and Jeffreys prior (blue curve) represented for fixed values of $(\mu, \sigma) = (0, 1)$. The black dashed lines represent Laplace distributions with scale parameter equal to $1/\lambda$, for $\lambda \in \{5, 10, 15\}$. Note that Laplace distributions $p_{\mathcal{L}}(\cdot | 1/\lambda)$ approximate well $p_{\text{PC}}(\cdot | \lambda)$ when $\lambda \geq 10$. Right panel: credible intervals at 95% for PC and Laplace priors, resp. $I_{\text{PC}}(0.95 | \lambda)$ and $I_{\text{Lap}}(0.95 | \lambda)$.

upper bound $\xi = 1$ which could have an undesirable influence in the posterior distribution. Thus, for the least informative case, a value of $\lambda = 1$ is sufficiently small as it does not favor values close to 0 nor those close to 1. In the case when 0 is favoured with a high λ and the true value of ξ differs from 0, the estimation may be altered compared to the uninformative case: see [Appendix C.5](#) for an analysis on simulated data. For the two other parameters, one can consider Jeffreys’ rule on (r, ν) in order to obtain a non-informative prior for (r, ν) while keeping invariance to reparameterization property. ξ is therefore considered *a priori* independent of (r, ν) . In view of the Fisher information matrix in Equation (9), we obtain $p_{\text{J}}(r, \nu) \propto 1/\nu$. Similarly to Jeffreys prior in [Section 3.1](#), the resulting prior is improper but the following proposition can be shown:

Proposition 4 *The prior defined as $p(r, \nu, \xi) \propto p_{\text{PC}}(\xi)p_{\text{J}}(r, \nu) \propto p_{\text{PC}}(\xi)/\nu$ for the Poisson process for extremes yields a proper posterior distribution.*

The proof, detailed in [Appendix B](#), relies on a result of [Northrop and Attalides \(2016\)](#). Note that this result still holds if $p_{\text{PC}}(\xi)$ is replaced by its approximation through the Laplace distribution.

4 Experiments

The benefits of the orthogonal reparameterization in the Poisson process model are illustrated here on simulations and a real environmental dataset. [Appendix C](#) contains additional experiments, notably using Hamiltonian Monte Carlo (HMC) instead of MCMC ([Appendix C.1](#)), under various maximum domains of attraction ([Appendix C.2](#)), in other models than the Poisson process model, namely the GPD and GEV ones ([Appendix C.3](#)), using the ratio-of-uniforms algorithm instead of MCMC ([Appendix C.4](#)) and finally with replications and comparison with maximum likelihood ([Appendix C.5](#)). All experiments are done using PyMC3 library ([Salvatier et al., 2016](#)), and the corresponding code is available online¹.

4.1 Simulations with the Poisson process model

Data generation We start by comparing the different parameterizations on exceedances generated with the Poisson process model described in [Section 1.1](#). For a given value of (μ, σ, ξ) and hyperparameters (u, m) , the data generation proceeds in two steps: first, a number of events n_u is simulated using a Poisson distribution with parameter $\Lambda(I_u)$ as defined in [Section 1.1](#). Then, for each point $i \in \{1, \dots, n_u\}$, the position x_i knowing that $x_i \in I_u$ is sampled from a GPD with parameters $(u, \tilde{\sigma}, \xi)$, with $\tilde{\sigma} = \sigma + \xi(u - \mu)$. An example with $(m, u, \mu, \sigma, \xi) = (40, 30, 50, 15, -0.25)$ is detailed here, leading to an expected number of observations $\Lambda(I_u) \approx 126$.

Experimental setup For MCMC hyper-parameters such as number of chains, burn-in period per chain or initialization, we keep the default values suggested in the PyMC3 library: four chains (which corresponds to our number of cores) with 1 000 iterations each, and a burn-in period of 1 000 iterations. In addition to these choices, this library offers the possibility to choose among different sampling methods, such as the traditional Metropolis–Hastings algorithm, but also more modern MCMC algorithms like Hamiltonian Monte Carlo (HMC, [Neal, 2011](#)), or the No-U-Turn sampler (NUTS, [Hoffman and Gelman, 2014](#)) which is the default choice in PyMC3. We choose to compare the different reparameterizations on Metropolis–Hastings draws (after burn-in), and the behaviour on NUTS is also investigated [Appendix C.1](#). We show that 1 000 iterations are sufficient for the chains to converge when the parameterization is well chosen. However, note that the algorithm only takes a few seconds to run, so this number of iterations can easily be increased for real data applications, as done in [Section 4.2](#). Finally, Jeffreys prior (computed in [Section 3.1](#)) is chosen for all configurations, but experiments have shown similar results with the PC prior of [Section 3.2](#).

Convergence diagnostic Our aim is to discriminate the different parameterizations according to the rate of convergence of the MCMC chains to their target. Different indicators exist to assess the quality of MCMC approximation. First, given a finite number of samples, autocorrelation plots as functions of lag measure how good the posterior approximation is, as the dependence between the chain elements reduces the effective information available for inference. To measure this, a common practice relies on the effective sample size, defined as $\text{ESS} = MN(1 + 2 \sum_{t=1}^{\infty} \rho_t)^{-1}$, with M the number of chains of size N , and ρ_t the autocorrelation at lag t . The ESS corresponds to

an equivalent number of independent draws, and so quantifies the amount of effective data for estimation (e.g., Gelman et al., 2013, Section 11.5). Here, the evolution of ESS with the number of draws is reported for each configuration. To complete the diagnostic, the potential scale reduction factor (commonly denoted by \hat{R}) also aims at bringing an indication about the state of convergence by computing the ratio of two estimators of the posterior variance. Generally $\hat{R} \geq 1$, and if it is greater than a given threshold, a convergence issue is raised. We use here a refinement of \hat{R} named \hat{R}_∞ (Moins et al., 2023), based on a local version $\hat{R}(x)$ which aims at ensuring the convergence at a given quantile x of the distribution. Then, \hat{R}_∞ is defined as the supremum of the $\hat{R}(x)$ values: $\hat{R}_\infty := \sup_{x \in \mathbb{R}} \hat{R}(x)$. This scalar summary amounts to considering the value of $\hat{R}(x)$ associated with the worse quantile approximation by the MCMC chains.

Results Results are reported in Figure 2, with four parameterizations that are compared for MCMC efficiency. (i) The orthogonal parameterization (r, ν, ξ) of Equation (7), and three triplets (μ, σ, ξ) associated with the following choices of m : (ii) the original m (same as the one used for generation), (iii) $m = n_u$ which is the choice of Wadsworth et al. (2010) and the package `revdbayes` (Northrop, 2023), and (iv) $m \in [m_1, m_2]$ as suggested by Sharkey and Tawn (2017) (see Section 2).

In order to compare the same quantities, all convergence diagnostics are computed with the original parameterization (μ, σ, ξ) and m , consequently after a transformation of the chains for the other parameterizations. Figure 2 confirms that the orthogonal parameterization behaves best in the case $\xi < 0$: the Markov chains have the lowest autocorrelations, the lowest value of $\hat{R}(x)$ for almost all x , and this parameterization is the only one that satisfies the recommendation of ESS lower than 400 for estimation (Gelman et al., 2013). Conversely, the two parameterizations that suggest a change for m seem to suffer from a lack of convergence, even more than the original parameterization. For the cases $\xi > 0$ and $\xi = 0$ detailed in Appendix C.2, the orthogonal parameterization is still best, but the behaviour of the three other parameterizations is reversed: the one with no change for m is the one with the largest convergence issues. Some intuitions about the behaviour of parameterizations that rely on changing m , in particular in the case $\xi < 0$, can be found in Appendix A. We also refer to Appendix C.3 for a study of the GPD and GEV cases. As a conclusion, the orthogonal parameterization is effective in the three maximum domains of attraction, for both Poisson process and GPD models.

4.2 Case study on river flow data

We apply our framework to 36 160 daily measurements of the Garonne river flow (France), from 1915 to 2013.

Preprocessing Before selecting a threshold and running a MCMC algorithm, some common preprocessing steps on daily environmental data are required: first because of seasonality, only the rainy season from December to May is considered, which reduces the number of observations to 18 043. The observations are also not independent; an autocorrelation plot suggests a three-day correlation in measurements. Therefore, clusters of exceedances of parameter $r = 3$ days are considered here, which means that two exceedances that occurred in less than three days are merged as one observation (the largest one in the cluster). Previous EDF studies (see for instance Albert, 2018) agree with traditional threshold elicitation methods (e.g., Coles, 2001, Chapter 4) to consider a threshold of $u = 2000 \text{ m}^3/\text{s}$ for estimation. In the end, $n_u = 182$ clusters of exceedances are obtained and represented in Figure 3. A temporal trend could be suspected there. A possible way

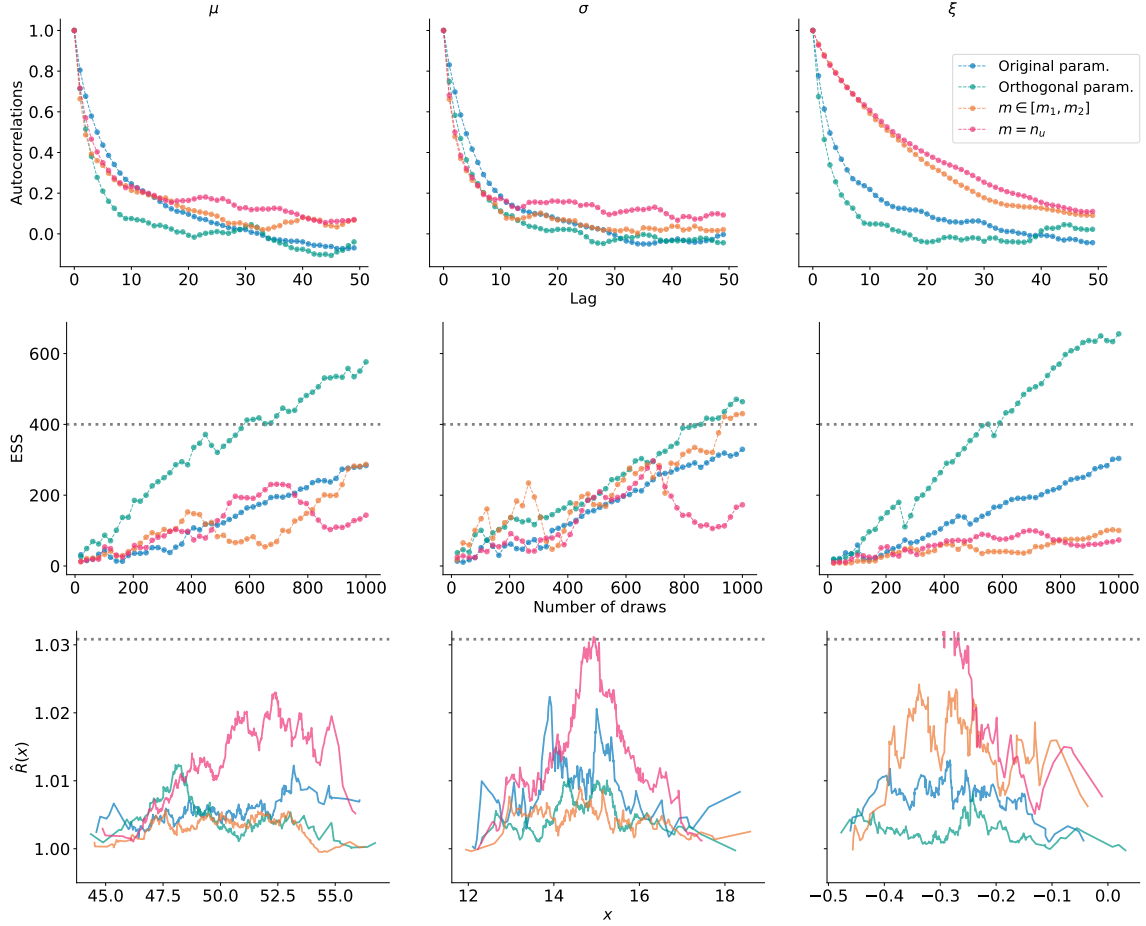


Figure 2: Convergence diagnostic plots for Poisson parameters (μ, σ, ξ) with $\xi < 0$, after 1000 Metropolis–Hastings draws and a burn-in of 1000, for four different parameterizations: the original one (in blue), the Sharkey and Tawn (2017) update with $m \in [\hat{m}_1, \hat{m}_2]$ (in orange), the Wadsworth et al. (2010) update with $m = n_u$ (in magenta), and the orthogonal parameterization (in green). Top row: autocorrelations as functions of the lag. Second row: evolution of ESS with the number of draws (the gray line corresponds to value of 400 recommended in Gelman et al. (2013)). Bottom row: $\hat{R}(x)$ as a function of the quantile x , with the adapted threshold of 1.03 (Moins et al., 2023). Some curves are truncated for visibility purposes, as they are taking much larger values than the threshold.

to model such a phenomenon would be to include covariates in the orthogonal parameters, see the last paragraph of Section 2.

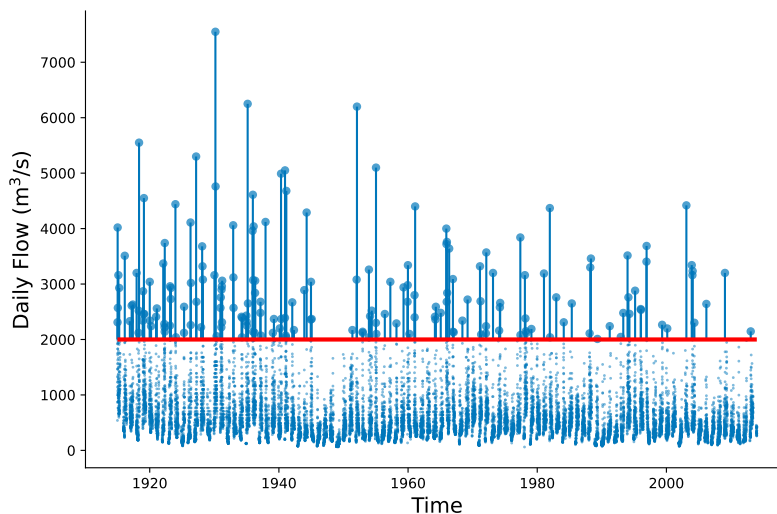


Figure 3: Plot of $n_u = 182$ exceedances of the Garonne river flow between 1915 and 2013 above the threshold $u = 2000$ (represented in red).

Return level estimation We are interested in estimating the T -year return level ℓ_T , which is exceeded on average once every T years. This is obtained by solving the equation $G(\ell_T | \mu, \sigma, \xi) = 1 - 1/T$, with G the GEV cdf defined in Equation (1):

$$\ell_T = \mu - \frac{\sigma}{\xi} \left(1 - (-\log(1 - 1/T))^{-\xi} \right). \quad (14)$$

Here, as the data span 99 years, we fix $m = 99$ in order to obtain parameters associated with annual maxima. The same setup as in Section 4.1 is then run with 5 000 draws from Metropolis–Hastings algorithm with the orthogonal parameterization. Convergence diagnostic values are reported in Figure 4 and show no evidence of lack of convergence, along with a very satisfactory effective sample size for estimation (final values can be found in Table 1 along with \hat{R}_∞ for each parameter). Results of posterior summaries for (μ, σ, ξ) are reported in Table 1: looking at the posterior for ξ , the three maximum domains of attraction cannot be excluded, although the 95% credible interval (CI) is tight around zero. This may suggest that $\xi = 0$ and an exponential decrease of the survival function. Return levels for annual maxima are displayed in the left panel of Figure 5, and show that the model seems to fit the data correctly. These curves are obtained by computing the mean and 2.5%/97.5% quantiles on the posterior distribution of ℓ_T for any given return period T . This is more accurate than the version where pointwise posterior quantities of (μ, σ, ξ) are plugged in Equation (14) (see Jonathan et al., 2021, for a comparison). The obtained posterior mean of ℓ_T , is 6 949 m³/s for the 100-year level and 9 266 m³/s for the 1 000-year one. These results corroborate a study conducted in Albert et al. (2020), where the estimated value of 10 000 m³/s for the 1 000-year return level belongs to the credible interval in Figure 5.

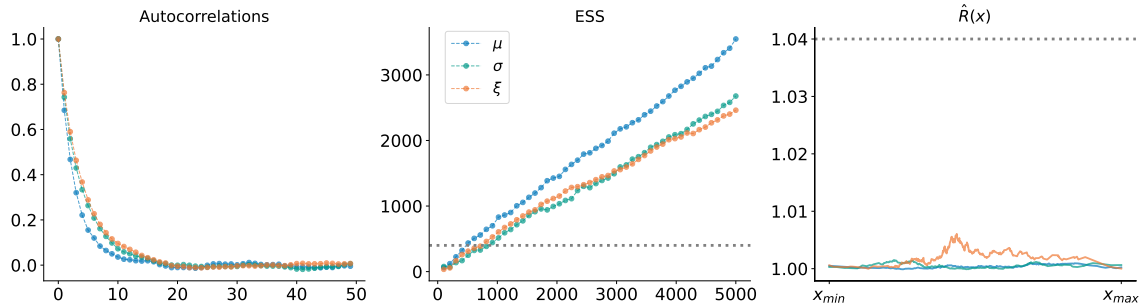


Figure 4: Convergence diagnostic plots for Garonne river flow data, after 5000 Metropolis–Hastings draws and a burn-in of 1000. Left: autocorrelations as functions of the lag. Middle: evolution of ESS with the number of draws (the gray line corresponds to value of 400 recommended in [Gelman et al. \(2013\)](#)). Right: $\hat{R}(x)$ as a function of the quantile x , with the adapted threshold of 1.04 (see [Moins et al., 2023](#)).

	Post. Mean	Post. SD	95%-CI	ESS	\hat{R}_∞
μ	2 560.8	84.1	[2 409.8, 2 724.1]	3 473	≈ 1.0
σ	919.6	73.2	[787.2, 1 063.3]	2 709	≈ 1.0
ξ	0.015	0.077	[-0.120, 0.164]	2 702	≈ 1.0

Table 1: Posterior summaries (mean, standard deviation (SD), credible interval (CI) at 95%) and convergence diagnostics (ESS and \hat{R}_∞) for (μ, σ, ξ) associated with annual maxima ($m = 99$).

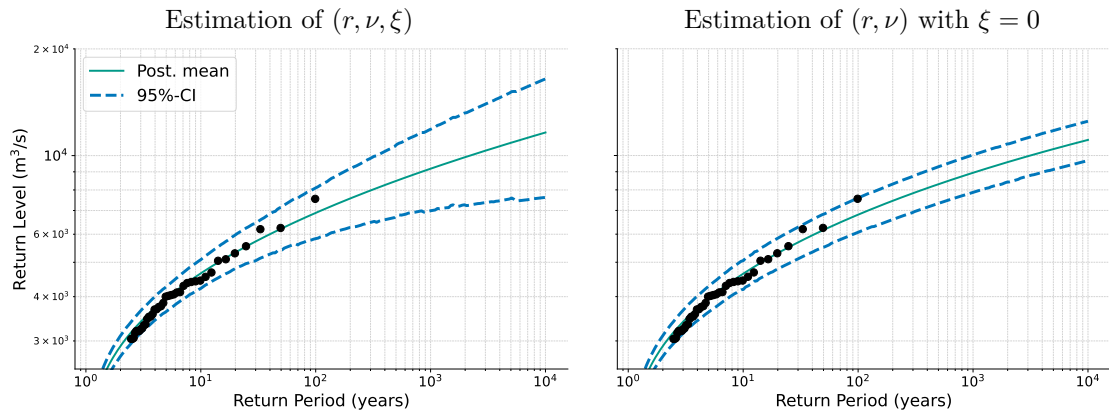


Figure 5: Return levels for annual maxima of Garonne flow data. Full green curves correspond to return levels obtained with posterior mean return level, and the dashed ones to the bounds of the 95% credible interval (CI). On the left, all three parameters (r, ν, ξ) are estimated, while on the right, only (r, ν) are estimated with the assumption that $\xi = 0$. Black points represent the observed annual maxima.

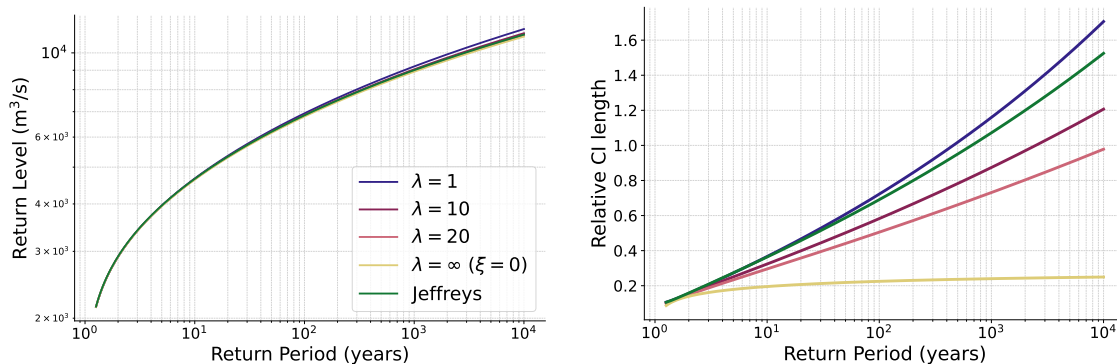


Figure 6: Comparison of return levels with different priors as functions of return period (log scale). On the left: return levels with posterior mean parameters. On the right: return level credible interval (CI) length relative to the point estimate (in %).

Prior influence on the return level estimation uncertainty Looking at the posterior distribution for ξ , one can reasonably make the assumption that $\xi = 0$ and therefore assume an exponential decrease for the survival function of the river flow. In this case, the remaining location parameter μ and scale parameter σ can be estimated with fixed $\xi = 0$. The resulting posterior summaries are very close to the ones of Table 1. As a result, the return level curves with posterior mean parameters (see Figure 5) are very similar in both cases. However, as the uncertainty on the shape parameter is excluded when fixing $\xi = 0$, the return levels credible intervals change drastically and become very concentrated around means, as shown in the right panel of Figure 5. In fact, this reflects that most of the uncertainty on the estimated return level is due to the estimation of the shape parameter, and so knowing its value greatly facilitates the extrapolation. PC priors allow us to navigate between these two extreme cases thanks to the hyperparameter λ . Looking at the left panel of Figure 6, it appears that the return level curves associated with posterior means are not affected by those differences of priors. However, the larger λ , the more information is added about the closeness of ξ to zero, and the smaller the length of the credible interval (note however that this does not give any guarantee on the estimation bias). This behaviour is illustrated on the right panel of Figure 6: denoting by $\ell_T^{(m)}$, $\ell_T^{(2.5\%)}$, and $\ell_T^{(97.5\%)}$ respectively the posterior mean, and the posterior quantiles at 2.5% and 97.5% of the return level, then the right plot in Figure 6 displays the length of the credible interval for the return level estimation, relatively to the estimator $\ell_T^{(m)}$: $(\ell_T^{(97.5\%)} - \ell_T^{(2.5\%)})/\ell_T^{(m)}$. This ratio is expected to grow with T , as the uncertainty increases in the tail. When $\lambda = 1$, this growth is similar to the one associated with Jeffreys prior, which can be seen as a noninformative case. For example, one can see that the size of the credible interval is already greater than the posterior estimation for the 1000-year return level (ratio greater than one). Using $\lambda = 10$ corresponds to a confidence of 95% of having ξ between -0.3 and 0.3 with the version approximated by a Laplace distribution (see the table in Figure 1), reduces by approximately 20% the size of the credible interval for $T = 1000$. The length when ξ is fixed at zero is drastically lower than in the other cases, even those concerning PC priors with large λ values.

5 Conclusion

In this paper we demonstrate the benefits of using an orthogonal parameterization in the sense of [Jeffreys \(1961\)](#) for Bayesian inference of extreme value models. First, orthogonal parameters facilitate the convergence of MCMC algorithms such as Metropolis–Hastings or NUTS (Section 2 and [Appendix A](#)). This improvement is “free” in the sense that it is obtained at no extra computational cost, except a simple change of variable if one interest lies in the original parameters (μ, σ, ξ) . This conclusion is confirmed by convergence diagnostics such as autocorrelation, effective sample size, and local \hat{R} , on simulations in the three maximum domains of attraction (Section 4.1 and [Appendix C](#)).

Secondly, the orthogonal parameterization also facilitates the computation of Jeffreys prior (Section 3.1): we show that this uninformative prior is defined for $\xi > -1/2$ and is improper, but leads to a proper posterior. Posterior propriety is a necessary condition for using this prior in practice when no external information is available. However, this uninformative case is actually far from the reality of most of the applications: even without any expert information, a shape parameter in the range $(-1, 1)$ already includes a vast majority of the distributions arising in natural phenomena. Therefore as an alternative, a PC prior on ξ can be used instead and allows users to control the prior knowledge they want to include on ξ (Section 3.2). In particular, it penalizes the values of ξ that move away from 0, and navigate between the uninformative case and the deterministic one where $\xi = 0$. In addition to its flexibility, this prior enjoys the same advantages as Jeffreys prior: invariance to reparameterization and posterior propriety. Additionally, it can be defined without any restriction for ξ if one uses the approximation by a Laplace distribution (otherwise, $\xi < 1$). This prior information on ξ impacts the posterior uncertainty around the return level estimation. By applying our framework on river flow data (Section 4.2), we showed that the length of the credible interval for the return level can be significantly reduced by adding prior information of ξ . However, the uncertainty around the return level can be quantified differently, by using the quantiles of the posterior predictive distribution defined in (6), see [Fawcett and Green \(2018\)](#) for a comparison. In future work, it would be interesting to also investigate the influence of the prior on the posterior predictive return levels.

Acknowledgement

We would like to thank the anonymous reviewers and an Editor for their careful reading and for providing us with valuable comments that helped us improving the manuscript. S. Girard acknowledges the support of the Chair Stress Test, Risk Management and Financial Steering, led by the École polytechnique and its Foundation and sponsored by BNP Paribas. J. Arbel acknowledges the support of the French National Research Agency (ANR-21-JSTM-0001).

References

- Albert, C. (2018). *Estimation des limites d’extrapolation par les lois de valeurs extrêmes. Application à des données environnementales*. PhD thesis (in French), Université Grenoble Alpes.
- Albert, C., Dutfoy, A., Gardes, L., and Girard, S. (2020). An extreme quantile estimator for the log-generalized Weibull-tail model. *Econometrics and Statistics*, 13:137–174.

- Belzile, L. R., Dutang, C., Northrop, P. J., and Opitz, T. (2022). A modeler’s guide to extreme value software.
- Betancourt, M. (2019). Incomplete Reparameterizations and Equivalent Metrics.
- Betancourt, M. and Girolami, M. (2015). Hamiltonian Monte Carlo for hierarchical models. In *Current trends in Bayesian methodology with applications*, pages 79–101. CRC Press.
- Bousquet, N. (2021). Bayesian Extreme Value Theory. In *Extreme Value Theory with Applications to Natural Hazards: From Statistical Theory to Industrial Practice*, pages 271–325. Springer, Cham.
- Browne, W. J., Steele, F., Golalizadeh, M., and Green, M. J. (2009). The use of simple reparameterizations to improve the efficiency of Markov chain Monte Carlo estimation for multilevel models with applications to discrete time survival models. *Journal of the Royal Statistical Society: Series A*, 172(3):579–598.
- Castellanos, M. E. and Cabras, S. (2007). A default Bayesian procedure for the generalized Pareto distribution. *Journal of Statistical Planning and Inference*, 137(2):473–483.
- Chavez-Demoulin, V. and Davison, A. C. (2005). Generalized additive modelling of sample extremes. *Journal of the Royal Statistical Society: Series C*, 54(1):207–222.
- Coles, S. (2001). *An Introduction to Statistical Modeling of Extreme Values*. Springer Series in Statistics. Springer-Verlag, London.
- Coles, S. G. and Powell, E. A. (1996). Bayesian methods in extreme value modelling: a review and new developments. *International Statistical Review*, 64(1):119–136.
- Cox, D. R. and Reid, N. (1987). Parameter orthogonality and approximate conditional inference. *Journal of the Royal Statistical Society: Series B*, 49(1):1–18.
- Diebolt, J., El-Aroui, M. A., Garrido, M., and Girard, S. (2005). Quasi-conjugate bayes estimates for GPD parameters and application to heavy tails modelling. *Extremes*, 8:57–78.
- Fawcett, L. and Green, A. C. (2018). Bayesian posterior predictive return levels for environmental extremes. *Stochastic Environmental Research and Risk Assessment*, 32(8):2233–2252.
- Gelfand, A. E., Sahu, S. K., and Carlin, B. P. (1995). Efficient parametrisations for normal linear mixed models. *Biometrika*, 82(3):479–488.
- Gelfand, A. E., Sahu, S. K., and Carlin, B. P. (1996). Efficient parametrizations for generalized linear mixed models. *Bayesian Statistics*, 5:48–74.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2013). *Bayesian Data Analysis*. CRC Press, New York, 3rd edition.
- Gelman, A., Vehtari, A., Simpson, D., Margossian, C. C., Carpenter, B., Yao, Y., Kennedy, L., Gabry, J., Bürkner, P.-C., and Modrák, M. (2020). Bayesian workflow.
- Gilks, W. R., Richardson, S., and Spiegelhalter, D. (1995). *Markov Chain Monte Carlo in Practice*. CRC press, Boca Raton.

- Gilleland, E. and Katz, R. W. (2016). extRemes 2.0: an extreme value analysis package in R. *Journal of Statistical Software*, 72:1–39.
- Haan, L. and Ferreira, A. (2006). *Extreme value theory: an introduction*. Springer, New York.
- Hoffman, M. D. and Gelman, A. (2014). No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(1):1593–1623.
- Huzurbazar, V. S. (1950). Probability distributions and orthogonal parameters. *Mathematical Proceedings of the Cambridge Philosophical Society*, 46(2):281–284.
- Jeffreys, H. (1946). An invariant form for the prior probability in estimation problems. *Proceedings of the Royal Society of London: Series A. Mathematical and Physical Sciences*, 186(1007):453–461.
- Jeffreys, H. (1961). *The Theory of Probability*. Oxford Univ. Press, Oxford, 3rd edition.
- Jóhannesson, Á. V., Siegert, S., Huser, R., Bakka, H., and Hrafnkelsson, B. (2022). Approximate Bayesian inference for analysis of spatiotemporal flood frequency data. *The Annals of Applied Statistics*, 16(2):905–935.
- Jonathan, P., Randell, D., Wadsworth, J., and Tawn, J. (2021). Uncertainties in return values from extreme value analysis of peaks over threshold using the generalised Pareto distribution. *Ocean Engineering*, 220:107725.
- Kotz, S. and Nadarajah, S. (2000). *Extreme Value Distributions: Theory and Applications*. Imperial College Press, London.
- Leadbetter, M., Lindgren, G., and Rootzén, H. (1983). *Extremes and Related Properties of Random Sequences and Processes*. Springer, New York.
- Moins, T., Arbel, J., Dutfoy, A., and Girard, S. (2023). On the use of a local \hat{R} to improve MCMC convergence diagnostic.
- Neal, R. M. (2011). MCMC using Hamiltonian dynamics. In *Handbook of Markov Chain Monte Carlo*, pages 113–162. Chapman and Hall/CRC.
- Northrop, P. J. (2022). *rust: Ratio-of-Uniforms Simulation with Transformation*. <https://paulnorthrop.github.io/rust/>, <https://github.com/paulnorthrop/rust>.
- Northrop, P. J. (2023). *revdbayes: Ratio-of-Uniforms Sampling for Bayesian Extreme Value Analysis*. <https://paulnorthrop.github.io/revdbayes/>, <https://github.com/paulnorthrop/revdbayes>.
- Northrop, P. J. and Attalides, N. (2016). Posterior propriety in Bayesian extreme value analyses using reference priors. *Statistica Sinica*, 26(2):721–743.
- Opitz, T., Huser, R., Bakka, H., and Rue, H. (2018). INLA goes extreme: Bayesian tail regression for the estimation of high spatio-temporal quantiles. *Extremes*, 21(3):441–462.
- Papaspiliopoulos, O., Roberts, G. O., and Sköld, M. (2003). Non-centered parameterisations for hierarchical models and data augmentation. *Bayesian Statistics*, 7:307–326.

- Pickands, J. (1975). Statistical inference using extreme order statistics. *Annals of Statistics*, 3(1):119–131.
- Robert, C. P. (2007). *The Bayesian Choice: from Decision-Theoretic Foundations to Computational Implementation*. Springer, New York.
- Roberts, G. O. and Polson, N. G. (1994). On the geometric convergence of the Gibbs sampler. *Journal of the Royal Statistical Society: Series B*, 56(2):377–384.
- Roberts, G. O. and Sahu, S. K. (1997). Updating schemes, correlation structure, blocking and parameterization for the Gibbs sampler. *Journal of the Royal Statistical Society: Series B*, 59:291–317.
- Salvatier, J., Wiecki, T. V., and Fonnesbeck, C. (2016). Probabilistic programming in Python using PyMC3. *PeerJ Computer Science*, 2:e55.
- Sharkey, P. and Tawn, J. A. (2017). A Poisson process reparameterisation for Bayesian inference for extremes. *Extremes*, 20(2):239–263.
- Simpson, D., Rue, H., Riebler, A., Martins, T. G., and Sørbye, S. H. (2017). Penalising model component complexity: A principled, practical approach to constructing priors. *Statistical Science*, 32(1):1–28.
- Stephenson, A. (2016). Bayesian Inference for Extreme Value Modelling. In *Extreme Value Modeling and Risk Analysis: Methods and Applications*, pages 257–280. Chapman & Hall/CRC: Boca Raton, Florida.
- Stephenson, A. and Tawn, J. (2004). Bayesian inference for extremes: Accounting for the three extremal types. *Extremes*, 7(4):291–307.
- Tibshirani, R. and Wasserman, L. (1994). Some aspects of the reparametrization of statistical models. *Canadian Journal of Statistics*, 22:163–173.
- Van der Vaart, A. W. (2000). *Asymptotic Statistics*. Cambridge University Press, Cambridge.
- Wadsworth, J. L., Tawn, J. A., and Jonathan, P. (2010). Accounting for choice of measurement scale in extreme value modeling. *Annals of Applied Statistics*, 4(3):1558–1578.
- Woutersen, T. (2011). Consistent estimation and orthogonality. *Advances in Econometrics*, 27(1):155–178.

Appendix A Approaching orthogonality by choosing $m = n_u$

Sharkey and Tawn (2017) suggests to take a value of the scaling factor m that minimises the off-diagonal terms of the asymptotic covariance matrix (that is the inverse Fisher information matrix), denoted by $\text{ACov} := \mathcal{I}^{-1}(\mu, \sigma, \xi)$. Those terms exist only if $\xi > -1/2$ (see Proposition 2

and its proof in [Appendix B](#)) and can be written as functions of $x = -\frac{1}{\xi} \log \left\{ 1 + \xi \left(\frac{u-\mu}{\sigma} \right) \right\}_+$, σ , and ξ as:

$$\begin{aligned} \text{ACov}_{\mu,\sigma} &= \frac{\sigma^2}{m\xi^2} e^x \left(\xi^3 + (1+\xi)(1+2\xi + \xi(1+\xi)x^2 - (1+3\xi)x + e^{-\xi x}(1+2\xi)(x-1)) \right), \\ \text{ACov}_{\mu,\xi} &= \frac{\sigma}{m\xi^2} e^x (1+\xi) \left(\xi(1+\xi)x - (1+2\xi)(1 - e^{-\xi x}) \right), \\ \text{ACov}_{\sigma,\xi} &= \frac{\sigma}{m} e^x (1+\xi) \left((1+\xi)x - 1 \right). \end{aligned}$$

Denote by $\rho_{\cdot,\cdot}$ the asymptotic correlation between two out of the three parameters, the authors note that a range of values may also work for m between m_1 and m_2 , where

$$m_1 = \underset{m}{\operatorname{argmin}} \{ |\rho_{\mu,\sigma}| + |\rho_{\mu,\xi}| \} \text{ and } m_2 = \underset{m}{\operatorname{argmin}} \{ |\rho_{\mu,\sigma}| + |\rho_{\sigma,\xi}| \}.$$

They also find on their experiments that m_1 cancels $\rho_{\mu,\sigma}$, and that m_2 cancels $\rho_{\sigma,\xi}$. A numerical method is used in [Sharkey and Tawn \(2017\)](#) to approximate m_1 and m_2 as functions of ξ . Therefore, this approach requires to study the roots x_1 of $\text{ACov}_{\sigma,\xi}$ and x_2 of $\text{ACov}_{\mu,\sigma}$ to respectively derive $\hat{m}_1(\xi)$ and $\hat{m}_2(\xi)$. Without any approximation, we directly have $x_1 = 1/(1+\xi)$ as the unique root for $\text{ACov}_{\sigma,\xi}$. Moreover, as $\xi > -1/2$, we have $x_1 > 0$, which motivates us to study the sign of the root x_2 for $\text{ACov}_{\mu,\sigma}$. Indeed, if x_2 is unique and $x_2 < 0$, then the choice $x = 0$ which cancels the third asymptotic covariance $\text{ACov}_{\mu,\xi}$ will always be reasonable as it will stay in the targeted interval, between the two other roots. In addition, $x = 0$ corresponds to the choice $m = r$ (which in practice translates into $m = n_u$), and is a simple choice as it does not require any estimation of ξ . The interest of the choice $m = n_u$ has already been mentioned in [Wadsworth et al. \(2010\)](#) to improve the mixing property of the chain. Unfortunately, a study of function $x \mapsto \text{ACov}_{\mu,\sigma}(x)$ shows that the properties of uniqueness and positivity of x_2 are only valid in the case where $\xi > 0$. In that case, the works of [Wadsworth et al. \(2010\)](#) and [Sharkey and Tawn \(2017\)](#) corroborate the choice of $m = n_u$. However, it is not the case anymore when $-1/2 < \xi < 0$. It can be shown that x_2 is not negative here, and worse, may not be unique. This can be seen as a counter-indication for frameworks that aim at reducing the three asymptotic covariances at the same time by tuning the scaling factor m .

Appendix B Proofs

Proof of Proposition 1 The log-likelihood l using the (r, ν, ξ) parameterization of Equation (7) can be written as:

$$\begin{aligned} l(r, \nu, \xi \mid \mathbf{x}, n_u) &= -r + n_u \log \left(\frac{r}{m} \right) - n_u \log(\nu) + n_u \log(1 + \xi) \\ &\quad - \left(1 + \frac{1}{\xi} \right) \sum_{i=1}^{n_u} \log \left\{ 1 + \frac{\xi(1+\xi)}{\nu} (x_i - u) \right\}_+. \end{aligned}$$

Under this form, one can directly see that r is orthogonal to ν and ξ . The second derivatives of l are

$$\begin{aligned}
\frac{\partial^2 l}{\partial r^2} &= -\frac{n_u}{r^2}, & \frac{\partial^2 l}{\partial r \partial \nu} &= 0, & \frac{\partial^2 l}{\partial r \partial \xi} &= 0, \\
\frac{\partial^2 l}{\partial \nu^2} &= \frac{n_u}{\nu^2} + \frac{\xi(1+\xi)^3}{\nu^4} \sum_{i=1}^{n_u} \frac{(x_i - u)^2}{\left\{1 + \frac{\xi(1+\xi)}{\nu}(x_i - u)\right\}_+^2} - \frac{2(1+\xi)^2}{\nu^3} \sum_{i=1}^{n_u} \frac{(x_i - u)}{\left\{1 + \frac{\xi(1+\xi)}{\nu}(x_i - u)\right\}_+}, \\
\frac{\partial^2 l}{\partial \nu \partial \xi} &= -\frac{(1+2\xi)(1+\xi)^2}{\nu^3} \sum_{i=1}^{n_u} \frac{(x_i - u)^2}{\left\{1 + \frac{\xi(1+\xi)}{\nu}(x_i - u)\right\}_+^2} + \frac{2(1+\xi)}{\nu^2} \sum_{i=1}^{n_u} \frac{(x_i - u)}{\left\{1 + \frac{\xi(1+\xi)}{\nu}(x_i - u)\right\}_+}, \\
\frac{\partial^2 l}{\partial \xi^2} &= -\frac{n_u}{(1+\xi)^2} + \frac{(1+2\xi)^2(1+\xi)}{\xi\nu^2} \sum_{i=1}^{n_u} \frac{(x_i - u)^2}{\left\{1 + \frac{\xi(1+\xi)}{\nu}(x_i - u)\right\}_+^2} \\
&\quad + \frac{2(1+\xi-\xi^2)}{\xi^2\nu} \sum_{i=1}^{n_u} \frac{(x_i - u)}{\left\{1 + \frac{\xi(1+\xi)}{\nu}(x_i - u)\right\}_+} - \frac{2}{\xi^3} \sum_{i=1}^{n_u} \log \left\{1 + \frac{\xi(1+\xi)}{\nu}(x_i - u)\right\}_+.
\end{aligned}$$

Focussing on the expectations, as we observe a Poisson process, the information is contained in the number n_u of observed points (we write N_u the corresponding random variable) and the position of jumping events x_i (we write X_i the corresponding random variable, with the same distribution as X). Here, N_u is distributed according to a Poisson distribution with parameter r , and $X - u$ is a GPD random variable with parameters $(\frac{\nu}{1+\xi}, \xi)$. For example, deriving the following expectations is the cornerstone to obtain the Fisher information matrix:

$$\begin{aligned}
\mathbb{E}_{N_u, X} \left[\sum_{i=1}^{N_u} \frac{(X_i - u)^2}{\left\{1 + \frac{\xi(1+\xi)}{\nu}(X_i - u)\right\}_+^2} \right] &= \mathbb{E}_{N_u} \left[\mathbb{E}_{X|N_u} \left[\sum_{i=1}^{N_u} \frac{(X_i - u)^2}{\left\{1 + \frac{\xi(1+\xi)}{\nu}(X_i - u)\right\}_+^2} \right] \right] \\
&= \mathbb{E}_{N_u} [N_u] \mathbb{E}_{X|N_u} \left[\frac{(X - u)^2}{\left\{1 + \frac{\xi(1+\xi)}{\nu}(X - u)\right\}_+^2} \right] \\
&= r \frac{1+\xi}{\nu} \int_u^{+\infty} (x - u)^2 \left\{1 + \frac{\xi(1+\xi)}{\nu}(x - u)\right\}_+^{-\frac{1}{\xi}-3} dx.
\end{aligned}$$

The above integral exists provided $\xi > -1/2$ and we obtain

$$\mathbb{E}_{N_u, X} \left[\sum_{i=1}^{N_u} \frac{(X_i - u)^2}{\left\{1 + \frac{\xi(1+\xi)}{\nu}(X_i - u)\right\}_+^2} \right] = \frac{2r\nu^2}{(1+\xi)^3(1+2\xi)}.$$

Similarly, the remaining expected values can be written as

$$\mathbb{E}_{N_u, X} \left[\sum_{i=1}^{N_u} \frac{(X_i - u)}{\left(1 + \frac{\xi(1+\xi)}{\nu}(X_i - u)\right)} \right] = \frac{r\nu}{(1+\xi)^2},$$

$$\mathbb{E}_{N_u, X} \left[\sum_{i=1}^{N_u} \log \left(1 + \frac{\xi(1+\xi)}{\nu}(X_i - u)\right) \right] = r\xi.$$

Plugging these values into the Fisher coefficients yields the result:

$$I(r, \nu, \xi) = \text{diag} \left(\frac{1}{r}, \frac{r}{\nu^2(1+2\xi)}, \frac{r}{(1+\xi)^2} \right).$$

Proof of Proposition 2 Let us show that the following integral exists for any $n_u \geq 1$:

$$C_{n_u} = \int_{\mathcal{S}} \frac{r^{1/2} e^{-r}}{\nu(1+\xi)(1+2\xi)^{1/2}} \left(\frac{r(1+\xi)}{m\nu} \right)^{n_u} \prod_{i=1}^{n_u} \left(1 + \frac{\xi(1+\xi)}{\nu}(x_i - u) \right)^{-1 - \frac{1}{\xi}} dr d\nu d\xi,$$

where \mathcal{S} is the integration domain:

$$\mathcal{S} = \left\{ (r, \nu, \xi) \in \mathbb{R}^3 \text{ s.t. } \xi > -\frac{1}{2}, r > 0, \nu \geq \{-\xi(1+\xi)((\max_i x_i) - u)\}_+ \right\}.$$

Let us consider the case of one observation ($n_u = 1$):

$$\begin{aligned} & \int_{-\frac{1}{2}}^{+\infty} (1+2\xi)^{-\frac{1}{2}} \int_0^{+\infty} r^{\frac{3}{2}} e^{-r} \int_{\{-\xi(1+\xi)(x-u)\}_+}^{+\infty} \nu^{-2} \left(1 + \frac{\xi(1+\xi)}{\nu}(x-u) \right)^{-\frac{1}{\xi}-1} d\nu dr d\xi \\ &= \int_{-\frac{1}{2}}^0 (1+2\xi)^{-\frac{1}{2}} \int_0^{+\infty} r^{\frac{3}{2}} e^{-r} \int_{-\xi(1+\xi)(x-u)}^{+\infty} \nu^{-2} \left(1 + \frac{\xi(1+\xi)}{\nu}(x-u) \right)^{-\frac{1}{\xi}-1} d\nu dr d\xi \\ &+ \int_0^{+\infty} (1+2\xi)^{-\frac{1}{2}} \int_0^{+\infty} r^{\frac{3}{2}} e^{-r} \int_0^{+\infty} \nu^{-2} \left(1 + \frac{\xi(1+\xi)}{\nu}(x-u) \right)^{-\frac{1}{\xi}-1} d\nu dr d\xi \\ &= \int_{-\frac{1}{2}}^0 (1+2\xi)^{-\frac{1}{2}} \int_0^{+\infty} r^{\frac{3}{2}} e^{-r} \left[\frac{1}{(1+\xi)(x-u)} \left(1 + \frac{\xi(1+\xi)}{\nu}(x-u) \right)^{-\frac{1}{\xi}} \right]_{-\xi(x-u)(\frac{r}{m})^\xi}^{+\infty} dr d\xi \\ &+ \int_0^{+\infty} (1+2\xi)^{-\frac{1}{2}} \int_0^{+\infty} r^{\frac{3}{2}} e^{-r} \left[\frac{1}{(1+\xi)(x-u)} \left(1 + \frac{\xi(1+\xi)}{\nu}(x-u) \right)^{-\frac{1}{\xi}} \right]_0^{+\infty} dr d\xi \\ &= \frac{1}{(x-u)} \int_{-\frac{1}{2}}^{+\infty} (1+\xi)^{-1} (1+2\xi)^{-\frac{1}{2}} \int_0^{+\infty} r^{\frac{3}{2}} e^{-r} dr d\xi \\ &= \frac{3\pi^{\frac{3}{2}}}{4(x-u)} < \infty. \end{aligned}$$

Therefore, the posterior is proper for $n_u = 1$. It is well-known that it stays so for $n_u > 1$ as can be seen by induction. For instance for $n_u = 2$, the posterior writes

$$p(\boldsymbol{\theta} \mid x_1, x_2) \propto p(x_1, x_2 \mid \boldsymbol{\theta}) p(\boldsymbol{\theta}) = p(x_2 \mid \boldsymbol{\theta}) p(x_1 \mid \boldsymbol{\theta}) p(\boldsymbol{\theta}) \propto p(x_2 \mid \boldsymbol{\theta}) p(\boldsymbol{\theta} \mid x_1) \leq p(\boldsymbol{\theta} \mid x_1)$$

which is integrable.

Proof of Proposition 4 Similarly to the proof of Proposition 2, the aim is to show the existence of the following integral for any n_u :

$$C_{n_u} = \int_{\mathcal{S}} \frac{p_{\text{PC}}(\xi | \lambda)}{\nu} e^{-r} \left(\frac{r}{m}\right)^{n_u} \left(\frac{\nu}{1+\xi}\right)^{-n_u} \prod_{i=1}^{n_u} \left(1 + \frac{\xi(1+\xi)}{\nu}(x_i - u)\right)^{-1-\frac{1}{\xi}} dr d\nu d\xi,$$

with $p_{\text{PC}}(\xi | \lambda)$ defined in Equation (13), and \mathcal{S} the following integration domain:

$$\mathcal{S} = \left\{ (r, \nu, \xi) \in \mathbb{R}^3 \text{ s.t. } \xi < 1, r > 0, \nu \geq \{-\xi(1+\xi)(\max_i x_i - u)\}_+ \right\}.$$

In the general case for n_u , we have

$$\begin{aligned} C_{n_u} &= \frac{\Gamma(n_u + 1)}{m^{n_u}} \int_{-\infty}^1 \int_{\{-\xi(1+\xi)(x-u)\}_+}^{+\infty} \frac{p_{\text{PC}}(\xi | \lambda)}{\nu} \left(\frac{\nu}{1+\xi}\right)^{-n_u} \\ &\quad \prod_{i=1}^{n_u} \left(1 + \frac{\xi(1+\xi)}{\nu}(x_i - u)\right)^{-1-\frac{1}{\xi}} d\nu d\xi \\ &= \frac{\Gamma(n_u + 1)}{m^{n_u}} \int_{-\infty}^1 \int_{\{-\xi(x-u)\}_+}^{+\infty} \frac{p_{\text{PC}}(\xi | \lambda)}{\sigma} \sigma^{-n_u} \prod_{i=1}^{n_u} \left(1 + \xi \left(\frac{x_i - u}{\sigma}\right)\right)^{-1-\frac{1}{\xi}} d\sigma d\xi. \end{aligned}$$

The remaining integral corresponds to the normalizing constant of the posterior distribution of a GPD model with a prior of the form $p(\sigma, \xi) \propto p(\xi)/\sigma$. Since $p(\xi)$ is a proper density, Theorem 1 in Northrop and Attalides (2016) allows us to conclude that C_{n_u} is finite for any $n_u \geq 1$. Note that this result remains true with $p_{\text{PC}}(\xi | \lambda)$ replaced by a Laplace distribution as suggested in Section 3.2, since the prior on ξ remains proper.

Appendix C Additional experiments

Appendix C.1 Simulations using an Hamiltonian Monte Carlo algorithm

Hamiltonian Monte Carlo (HMC, Neal, 2011) and its variants such as NUTS (Hoffman and Gelman, 2014) are MCMC methods with a Markov kernel based on trajectories of particles computed using Hamiltonian dynamics. As a consequence, the performance of these methods is also sensitive to the choice of the parameterization (see Betancourt, 2019 for a formalization of the problem). We performed the same experiments as those in Section 4.1 and Appendix C.2, using 500 NUTS iterations instead of 1000 Metropolis–Hastings draws. The results obtained here are similar, and show that the orthogonal parameterization improves the efficiency of NUTS sampling. The case $\xi > 0$ is illustrated in Figure C.7 with the same configuration as the one described in the first paragraph of Appendix C.2. We observe similar trends as those in Figure C.8: changing the value of m improves convergence, and using the orthogonal parameterization is even better. Moreover, NUTS seems to be more efficient on the three cases than with Metropolis–Hastings, as the chains seem to be less correlated compared to their equivalent in Figure C.8, and the ESS can even be greater than the number of draws.

Appendix C.2 Simulations in other maximum domains of attraction

We study the influence of parameterizations for MCMC convergence in cases where $\xi > 0$ and $\xi = 0$.

Example with $\xi > 0$ Here, we set $(m, u, \mu, \sigma, \xi) = (5, 10, 30, 15, 0.7)$, which leads to an expected number of observations $r \approx 239$. Looking at autocorrelations, ESS and $\hat{R}(x)$ curves in Figure C.8, we can first confirm the result of Sharkey and Tawn (2017) about the inefficiency of Metropolis–Hastings with the original parameterization: high autocorrelations, high $\hat{R}(x)$ (around 1.7 for the highest) and almost zero ESS even after 1 000 iterations indicate a severe convergence issue. Changing the value of m before the MCMC algorithm as suggested by Sharkey and Tawn (2017) or by Wadsworth et al. (2010) improves inference significantly. Still, our orthogonal parameterization is even more efficient, especially for the estimation of the tail parameter ξ : the autocorrelation reduces even more rapidly with the lag, and the ESS increases faster with the number of draws. With the recommendations of ESS ≥ 400 for estimation (Gelman et al., 2013), our experimental setup is satisfactory only in the orthogonal case because of ξ . In contrast, more iterations are required to fulfill this condition for the parameterization recommended by Sharkey and Tawn (2017).

Example with $\xi = 0$ Finally when $\xi = 0$, the GPD and therefore the intensity $\Lambda(I_u)$ of the Poisson process defined in Section 1.1 reduce to an exponential model with location and scale parameters. Figure C.9 shows an example in this case with $(m, u, \mu, \sigma, \xi) = (20, 20, 25, 5, 0)$, leading to $r \approx 54$ expected observations. Similarly to the case $\xi > 0$ in Section 4.1, this example illustrates that updating m like Sharkey and Tawn (2017) or Wadsworth et al. (2010) is beneficial for MCMC convergence, but less than using orthogonal parameterization. In the same way as in the two other maximum domains of attraction, this parameterization is the most efficient one for the convergence of Metropolis–Hastings algorithm.

Appendix C.3 GPD and GEV case

In the particular case of GPD (defined in Equation (2)) that arises in the traditional peaks over threshold model, the same observation can be made about the benefits of an orthogonal parameterization for (σ, ξ) . More precisely, the transformation $(\nu, \xi) = (\sigma(1 + \xi), \xi)$ leads to an orthogonal Fisher information matrix for GPD (Chavez-Demoulin and Davison, 2005), and improves MCMC convergence as shown in Figure C.10. The same experimental setup as in the Poisson process case is used here, with a choice of $(\sigma, \xi) = (5, -0.1)$ and $u = 25$. Again, all plots in Figure C.10 show that the chains mixing is satisfactory only in the case of an orthogonal parameterization, while the original parameterization requires more iterations to be effective for inference. Up to our knowledge, there is no orthogonal parameterization for the GEV likelihood known in the literature. However, it should be noted that the parameters of the Poisson process model (μ, σ, ξ) correspond to those of the block maxima framework with m blocks (see Section 1.1). Consequently, we should expect a similar convergence issue for parameters (μ, σ, ξ) with GEV likelihood, and therefore an improvement in the MCMC convergence with the use of the orthogonal parameterization (r, ν, ξ) of the Poisson model.

Appendix C.4 Ratio-of-uniforms

The benefits of reparameterization for Bayesian inference can be extended to other sampling methods than MCMC. Typically, the efficiency of acceptance-rejection algorithms can be altered if the geometry of the acceptance region is too complex, and this can be due to correlation between parameters. The `rust` package (Northrop, 2022) implements such an acceptance-rejection algorithm dedicated to extreme value models, the so-called generalized ratio-of-uniforms method. It consists in simulating uniformly values in a region that encloses an acceptance region, where the ratio of

the obtained samples is distributed according to the target distribution (see [Gilks et al. \(1995, Chapter 5\)](#) for more details). As explained in the `revdbayes` documentation ([Northrop, 2023](#)) which is built upon `rust`, the efficiency of this method highly depends on the probability of acceptance p_a . `revdbayes` already includes the possibility to use the reparameterization suggested by [Wadsworth et al. \(2010\)](#) with $m = n_u$ for the Poisson process, along with a rotate option to reduce dependence. We add the orthogonal parameterization in the comparison, and show the results in [Figure C.11](#). We set $(m, u, \mu, \sigma, \xi) = (100, 0, 1, 1, -0.1)$, and draw 10 000 samples for three configurations. As expected, the orthogonal parameterization slightly improves the probability of acceptance compared to the case where $m = n_u = 110$, which is already significantly better than the case where m is not changed ($m = 100$). Note that this package operates a transformation of variable before sampling to improve normality, which in view of the bottom row in [Figure C.11](#), may not be necessary for the orthogonal parameters.

Appendix C.5 Replications and comparison with maximum likelihood

Despite the fact that the Bayesian paradigm comes with several benefits (briefly described in [Section 1.1](#)), one can be interested in the comparison with frequentist estimator such as maximum likelihood estimation (MLE). From a frequentist point of view, this involves extracting a pointwise estimator from the posterior distribution, such as the posterior mean, and replicate the experiment to estimate the mean squared error (MSE). The two steps of the Bayesian workflow we study here are expected to impact the performance of these estimators. A parameterization which leads to poor convergence of the MCMC chains will affect the accuracy of estimation, and the prior can add a bias that may or may not be advantageous to the estimation.

For different values of ξ_0 between -0.5 and 1 , we replicate 100 times the following experiment (this range includes a large number of models and allows us to have both Jeffreys and PC priors well defined): for $i = 1, \dots, 100$, we generate samples \mathbf{x}_i according to a Poisson process distribution with parameters $(m, u, \sigma, \xi) = (1, 10, 15, \xi_0)$ and μ in a way such that the expected number of points is equal to $r = 100$:

$$\mu = u - \frac{\sigma}{\xi_0}(100^{-\xi_0} - 1).$$

Then, we run MCMC chains with the same configuration as in [Section 4](#) and compute the posterior mean $\hat{\xi}_i = \mathbb{E}[\xi | \mathbf{x}_i]$. We these 100 experiments, we compute the MSE:

$$\text{MSE}(\xi_0) = \frac{1}{100} \sum_{i=1}^{100} (\hat{\xi}_i - \xi_0)^2.$$

First, we compare the different parameterizations for the Poisson process with the same Jeffreys prior. Results are displayed in the left panel of [Figure C.12](#), and illustrate the inaccuracy of the frameworks without reparameterization and with the update of [Sharkey and Tawn \(2017\)](#), due to lack of convergence of MCMC. This issue is getting worse as ξ_0 increases, and a bias/variance decomposition of the MSE shows that it is mostly due to the variance term. Then, for the same orthogonal parameterization, we compare Jeffreys prior, PC prior with a choice of $\lambda = 10$, and the MLE for the Poisson process, implemented using the `extRemes` package ([Gilleland and Katz, 2016](#)). Results in the right panel of [Figure C.12](#) show that the performance of the posterior mean estimation with Jeffreys prior is approximately the same as the MLE, except when ξ_0 is close to $-1/2$ where the asymptote behaviour of Jeffreys favours the estimation. This shows that, despite the

uninformative construction, this prior can favour negative values of ξ_0 close to $-1/2$. The behaviour of PC prior is, as expected, penalizing the values of ξ far from $\xi_0 = 0$. When ξ_0 is around zero, this prior outperforms Jeffreys' one and MLE, but assuming a value near zero when $|\xi_0|$ is large can add a large bias.

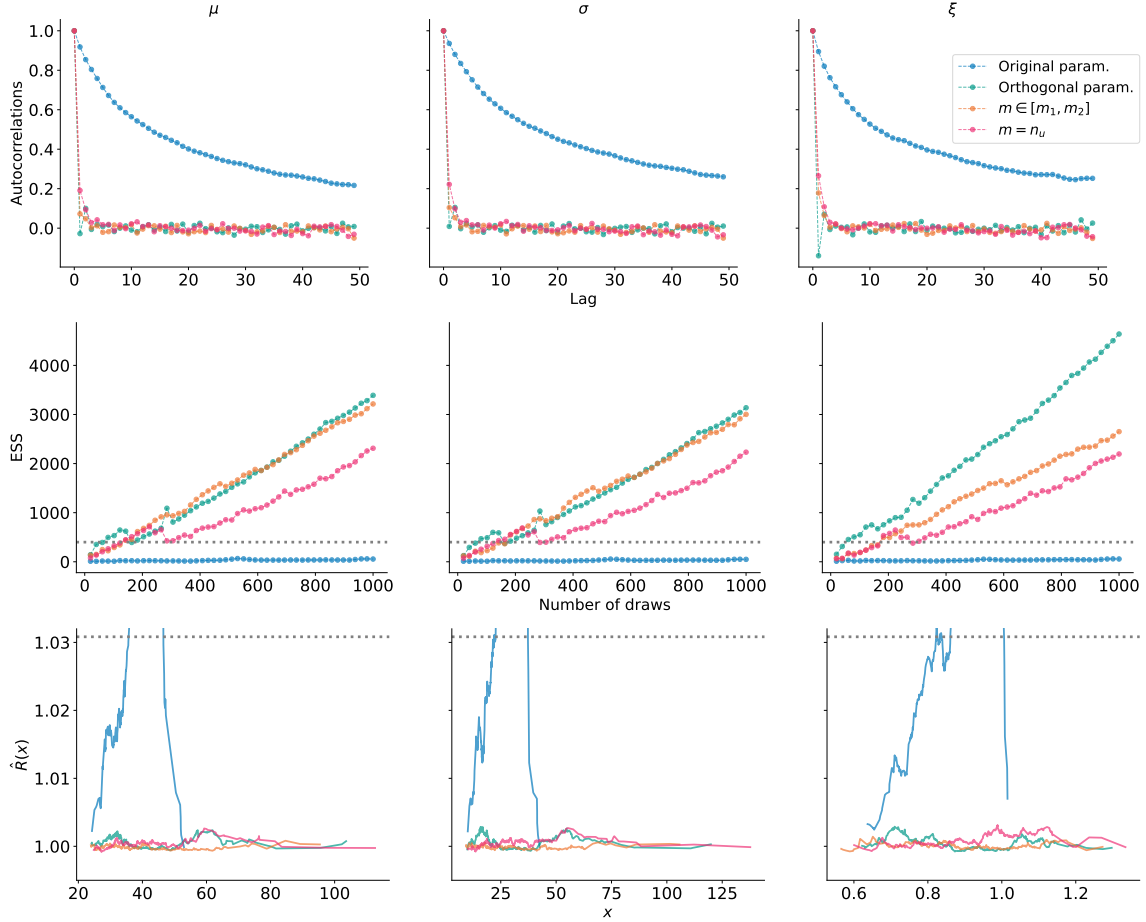


Figure C.7: Convergence diagnostic plots for Poisson parameters (μ, σ, ξ) with $\xi > 0$, after 500 NUTS draws and a burn-in of 1000, for four different parameterizations: the original one (in blue), the Sharkey and Tawn (2017) update with $m \in [\hat{m}_1, \hat{m}_2]$ (in orange), the Wadsworth et al. (2010) update with $m = n_u$ (in magenta), and the orthogonal parameterization (in green). Top row: autocorrelations as functions of the lag. Second row: evolution of ESS with the number of draws (the gray line corresponds to value of 400 recommended in Gelman et al. (2013)). Bottom row: $\hat{R}(x)$ as a function of the quantile x , with the adapted threshold of 1.03 (see Moins et al., 2023). The red curve is truncated for visibility purposes, as it is taking much larger values than the threshold.

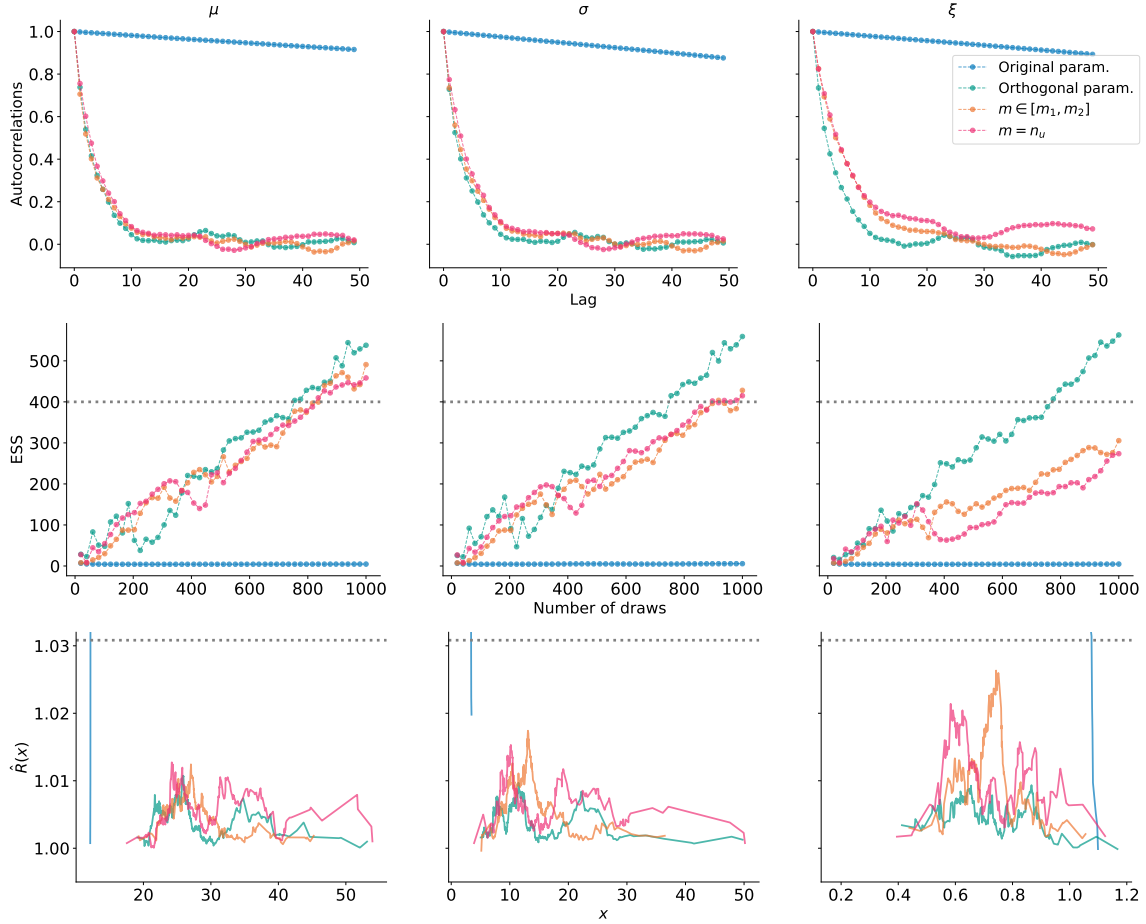


Figure C.8: Convergence diagnostic plots for Poisson parameters (μ, σ, ξ) with $\xi > 0$, after 1 000 Metropolis–Hastings draws and a burn-in of 1 000, for four different parameterizations: the original one (in blue), the Sharkey and Tawn (2017) update with $m \in [\hat{m}_1, \hat{m}_2]$ (in orange), the Wadsworth et al. (2010) update with $m = n_u$ (in magenta), and the orthogonal parameterization (in green). Top row: autocorrelations as functions of the lag. Second row: evolution of ESS with the number of draws (the gray line corresponds to value of 400 recommended in Gelman et al. (2013)). Bottom row: $\hat{R}(x)$ as a function of the quantile x , with the adapted threshold of 1.03 (see Moins et al., 2023). The red curve is truncated for visibility purposes, as it is taking much higher values than the threshold.

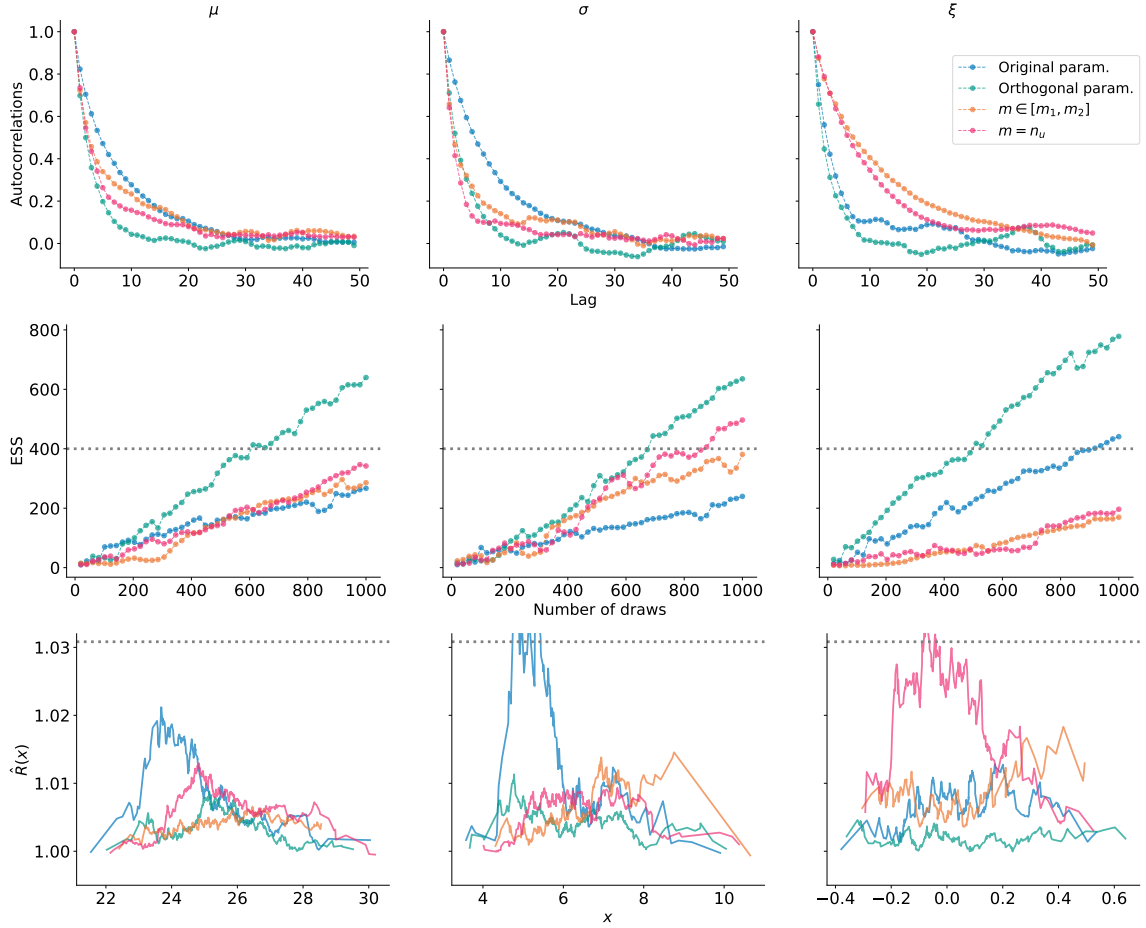


Figure C.9: Convergence diagnostic plots for Poisson parameters (μ, σ, ξ) with $\xi = 0$, after 1000 Metropolis–Hastings draws and a burn-in of 1000, for four different parameterizations: the original one (in blue), the Sharkey and Tawn (2017) update with $m \in [\hat{m}_1, \hat{m}_2]$ (in orange), the Wadsworth et al. (2010) update with $m = n_u$ (in magenta), and the orthogonal parameterization (in green). Top row: autocorrelations as functions of the lag. Second row: evolution of ESS with the number of draws (the gray line corresponds to value of 400 recommended in Gelman et al. (2013)). Bottom row: $\hat{R}(x)$ as a function of the quantile x , with the adapted threshold of 1.03 (see Moins et al., 2023).

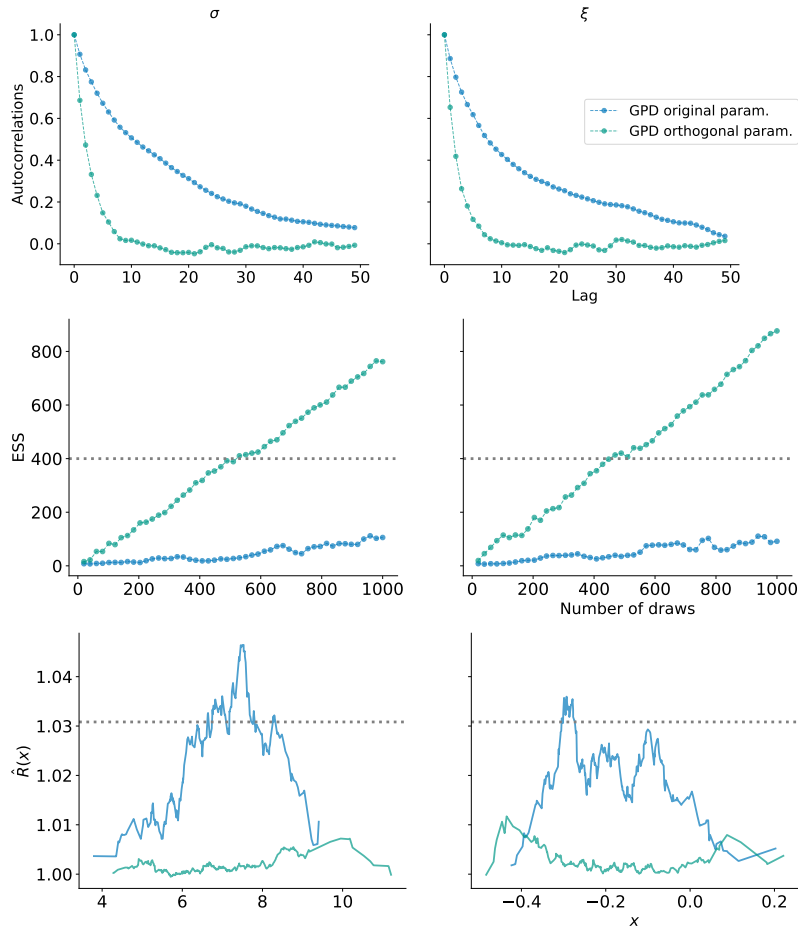
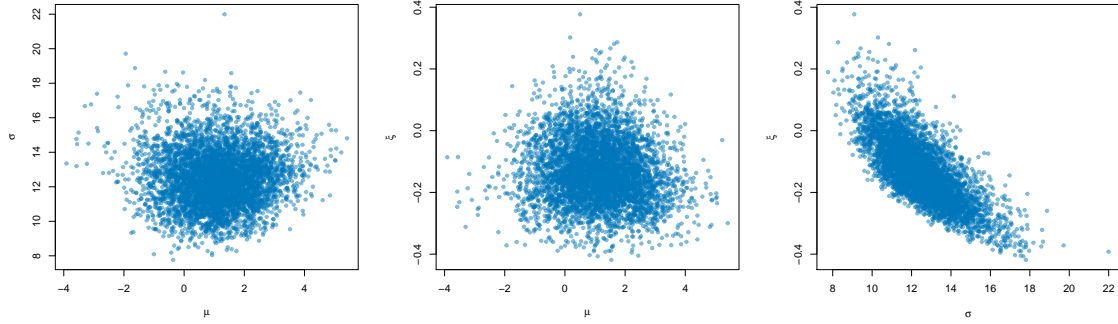
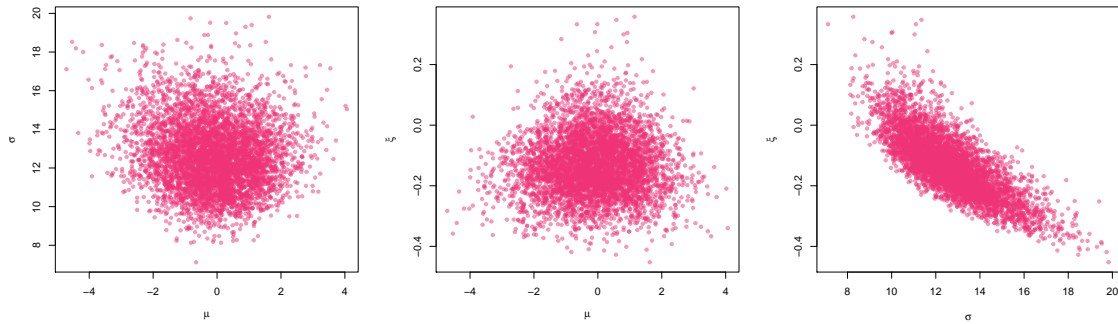


Figure C.10: Convergence diagnostic plots for GPD parameters (σ, ξ) with $\xi < 0$, after 1 000 Metropolis–Hastings draws and a burn-in of 1 000, for two parameterizations, the original (in blue) and the orthogonal one (in green). Top row: autocorrelations as functions of the lag. Second row: evolution of ESS with the number of draws (the gray line corresponds to value of 400 recommended in Gelman et al. (2013)). Bottom row: $\hat{R}(x)$ as a function of the quantile x , with the adapted threshold of 1.03 (see Moins et al., 2023).

Original parameters (μ, σ, ξ) ($p_a \approx 0.195$)



Original parameters (μ, σ, ξ) with $m = n_u$ ($p_a \approx 0.284$)



Orthogonal parameters (r, ν, ξ) ($p_a \approx 0.301$)

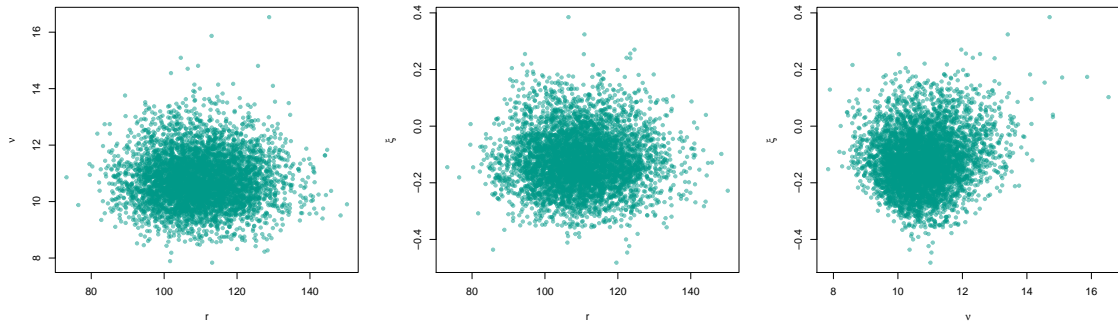


Figure C.11: Pairwise plots of parameter values simulated using the ratio of uniform method, for three parameterizations: the original one (in blue), the [Wadsworth et al. \(2010\)](#) update with $m = n_u$ (in magenta), and the orthogonal parameterization (in green). The probability of acceptance p_a reflects the efficiency of the sampling method.

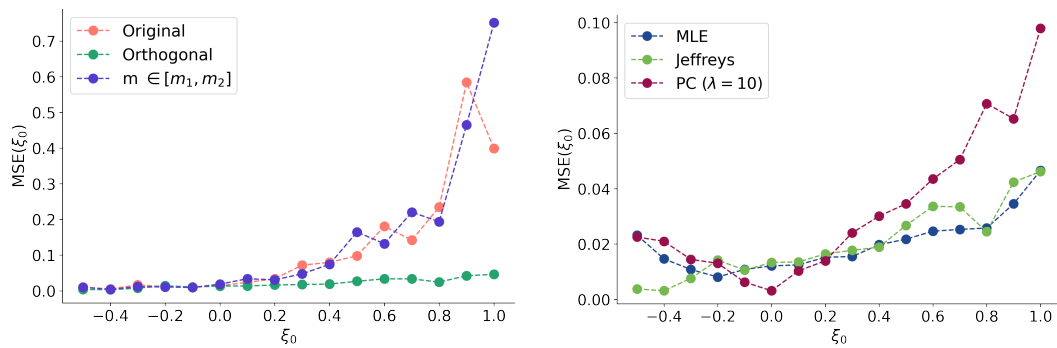


Figure C.12: Mean squared error (MSE) on the estimation of ξ for a true value $\xi_0 \in [-1/2, 1]$. The computation is done on 100 replications for each value of ξ_0 . Left panel: different parameterizations under Jeffreys prior. Right panel: Jeffreys and PC priors under orthogonal parameterization, along with MLE.