

Non-stationary conditional extremes of northern North Sea storm characteristics

P. Jonathan^{a*}, K. Ewans^b and D. Randell^a

Summary:

Characterising the joint structure of extremes of environmental variables is important for improved understanding of those environments. Yet many applications of multivariate extreme value analysis adopt models that assume a particular form of extremal dependence between variables without justification, or restrict attention to regions in which all variables are extreme. The conditional extremes model of Heffernan and Tawn (2004) provides one approach to avoiding these particular restrictions. Extremal marginal and dependence characteristics of environmental variables typically vary with covariates. Reliable descriptions of extreme environments should also therefore characterise any non-stationarity. Jonathan et al. (2013) extends the conditional extremes model of Heffernan and Tawn to include covariate effects, using Fourier representations of model parameters for single periodic covariates.

Here, we extend the work of Jonathan et al. (2013), introducing a general-purpose spline representation for model parameters as functions of multidimensional covariates, common to all inference steps. We use a non-crossing quantile regression to estimate appropriate non-stationary marginal quantiles simultaneously as functions of covariate; these are necessary as thresholds for extreme value modelling, and for standardisation of marginal distributions prior to application of the conditional extremes model. Then we perform marginal extreme value and conditional extremes modelling within a roughness-penalised likelihood framework, with cross-validation to estimate suitable model parameter roughness. Finally, we use a bootstrap re-sampling procedure, encompassing all inference steps, to quantify uncertainties in, and dependence structure of, parameter estimates and estimates of conditional extremes of one variate given large values of another.

We validate the approach using simulations from known joint distributions, the extremal dependence structures of which change with covariate. We apply the approach to joint modelling of storm peak significant wave height and associated storm peak period for extra-tropical storms at a northern North Sea location, with storm direction as covariate. We evaluate the impact of incorporating directional effects on estimates for conditional return values.

Keywords: non-stationarity; conditional extremes; covariate; spline;
non-crossing quantile regression; cross-validation; bootstrap.

^aShell Projects & Technology, Manchester, UK

^bSarawak Shell Bhd., 50450 Kuala Lumpur, Malaysia

*Correspondence to: P. Jonathan, Shell Projects & Technology, Brabazon House, Concord Business Park, Manchester M22 0RR, UK. E-mail: philip.jonathan@shell.com

1. INTRODUCTION

Characteristics of extreme environments are non-stationary in general; they vary with respect to covariates. Following the work of Davison and Smith (1990), numerous models for non-stationary extremes have been proposed. For example, Chavez-Demoulin and Davison (2005) describes smooth non-stationary generalized additive modelling for sample extremes, in which spline smoothers are incorporated into models for exceedances over high thresholds. Cooley et al. (2006) uses probabilistic inversion to infer the ages of climatic events based on observations of extreme lichen diameters. Eastoe (2009) uses a pre-processing approach to accommodate non-stationarity as part of a hierarchical model for multivariate extremes with covariates. Atyeo and Walshaw (2009) develop a region-based hierarchical model for extreme rainfall over the UK, incorporating spatial dependence and temporal trend. Gyarmati-Szabo et al. (2011) models threshold exceedances of air pollution concentrations via a non-homogeneous Poisson process with multiple change-points using reversible jump Markov chain Monte Carlo. Grigg and Tawn (2012) uses a censored generalised extreme value model for non-stationary threshold exceedances of extreme river flows. In oceanography and ocean engineering, marginal extremes of significant wave height and spectral peak period vary with wave direction and season (as demonstrated by, for example, Ewans and Jonathan 2008 and Jonathan and Ewans 2011). Coles and Walshaw (1994) develops a directional model for extreme wind speed. Coles and Casson (1998) proposes an extreme value model for hurricane wind speed with spatially-varying parameters. Butler et al. (2007) characterises changes in the occurrence and severity of storm surge events in the southern and central North Sea over the period 1955-2000. Mendez et al. (2008) introduces a seasonal peaks over threshold model for significant wave height. Ruggiero et al. (2010) examines increasing wave heights and extreme value projections for the U.S. Pacific Northwest. Kyselý et al. (2010) estimates extremes in climate change simulations using the peaks over threshold method with non-stationary threshold. Northrop and Jonathan (2011) discusses threshold

modelling of spatially dependent non-stationary extremes with application to hurricane-induced significant wave height. The characteristics of spectral peak period given large values of significant wave height also exhibit directional dependence in general (see, for example Jonathan et al. (2013)).

The conditional extremes model of Heffernan and Tawn (2004) provides a straightforward framework for estimating multivariate extremal dependence in the absence of covariates. The approach uses an asymptotic argument which conditions on one component of a random vector and finds the limiting conditional distribution of the remaining components as the conditioning variable becomes large. Conditions for the asymptotic argument to hold have been explored by Heffernan and Resnick (2007). To estimate the conditional extremes model for bivariate extremes of random variables \dot{X}_1 and \dot{X}_2 , and to simulate realisations under the model, the following procedure is appropriate. (a) Select a range of appropriate thresholds for threshold exceedance modelling for each variable in turn. (b) Estimate marginal generalised Pareto models for threshold exceedances for each variable in the sample in turn for different threshold choices, plot the values of parameter estimates as a function of threshold, and select the lowest threshold value per variable corresponding to appropriate parameter behaviour. (c) Transform \dot{X}_1 and \dot{X}_2 in turn to the Gumbel scale (to X_1 and X_2) using the probability integral transform. (d) Estimate conditional extremes models for $X_2|X_1$ (and $X_1|X_2$) for various choices of threshold of the conditioning variate, retain the estimated model parameters and residuals, plot the values of model parameter estimates and examine residuals as a function of threshold, and select the lowest threshold per variable consistent with modelling assumptions. (e) Simulate joint extremes on the standard Gumbel scale under the model, and transform realisations to the original scale using the probability integral transform.

In non-stationary marginal extreme value modelling of threshold exceedances, in addition to non-stationary forms for generalised Pareto scale (in particular) and shape parameters,

there is a growing literature favouring the adoption of non-stationary extreme value threshold (see, for example, Kysely et al. 2010, Thompson et al. 2010, Northrop and Jonathan 2011). Anderson et al. (2001) notes that the combination of non-stationary threshold and stationary generalised Pareto shape and scale is sufficient for modelling a sample of significant wave heights in the North Sea. Physical and statistical intuition suggest, when considering a non-stationary conditional extremes model for a given application, that non-stationary estimates should be sought at each of the modelling stages (a), (b) and (d) above in order, and adopted if justified statistically as in Section 5 below. Note however that physically plausible oceanographic examples corresponding, for instance, to stationary marginal but non-stationary conditional distributions are also conceivable.

A number of applications and refinements of the conditional extremes approach have been reported. Keef et al. (2009) examines spatial risk assessment for extreme river flows using the conditional extremes model. Keef et al. (2013a) proposes a variant of the conditional extremes model in which marginal transformation to Laplace rather than Gumbel scale is performed. The former has exponential tails on both sides and symmetry, capturing the exponential upper tail of the Gumbel required for modelling positive dependence but the symmetry also allows for negatively associated variables to be incorporated into the model parsimoniously. Keef et al. (2013b) proposes additional constraints within the conditional extremes model formulation, particularly relevant for negatively associated variables. Gilleland et al. (2013) use the conditional extremes model to estimate joint extremes of large-scale indicators for severe weather.

Estimation of joint extremal behaviour is of considerable interest in ocean engineering in particular (see, for example Ewans and Jonathan 2014). Jonathan et al. (2010) applies the conditional extremes model to storm peak significant wave height and associated spectral peak period for hindcast samples from different ocean basins. Jonathan et al. (2012) considers the joint modelling of directional ocean currents. There are numerous other approaches,

including the popular empirical method of Haver and Nyhus (1986), and the first order reliability method (FORM) and inverse-FORM (I-FORM) methodologies (for example Winterstein et al. 1993 and Winterstein et al. 1999), both of which attempt to characterise the marginal and conditional distributions of oceanographic variables. Recently, Jonathan et al. (2013) extended the conditional extremes model to incorporate the effect of covariates. The objective of the current work is to develop and evaluate an extended non-stationary conditional extremes model using common general-purpose penalised spline representations of model parameters with respect to multidimensional covariates. We illustrate the approach in application to joint inference for storm peak significant wave height (H_S) and associated storm peak period (T_P) with respect to storm direction as covariate.

For those less familiar with the physics of ocean waves and wave-structure interactions, significant wave height can be defined as four times the standard deviation of the ocean surface elevation for a time interval, and peak period as the reciprocal of the frequency at which the maximum wave energy is transmitted in the same interval. For wind-generated waves, the mean value of peak period increases with increasing significant wave height. Together, significant wave height (measured in metres) and peak period (measured in seconds) provide a useful low-dimensional summary of the ocean wave (energy) spectrum and therefore the characteristics of ocean waves for the time interval of interest. Kinsman (1984) and Stewart (2008) provide excellent introductions. Reliable long term measurements of significant wave height and associated wave characteristics for a specific ocean location of interest over a sufficiently long period of time are rarely available. Instead, hindcast time-series of ocean wave characteristics are generated using a physical model for the ocean environment, incorporating wind field and wind wave generation models in particular. The hindcast model is calibrated to observations of the environment from instrumented offshore facilities, moored buoys and satellite altimeters in the neighbourhood of the location (see, e.g. Cox and Swail 2000, Swail and Cox 2000). Storm peak significant wave height and

associated storm peak period are isolated from hindcast time-series of significant wave height and peak period corresponding to sea states of 3 hours duration for the whole historical period of interest, using the procedure described in Ewans and Jonathan (2008). Briefly, contiguous intervals of significant wave height above a low threshold are identified, each interval corresponding to a storm event. The maximum of significant wave height during the interval is termed the storm peak significant wave height for the storm. The value of peak period at the time of the storm peak significant wave height is referred to as the associated storm peak period. Storm peak significant wave height and associated peak period provide a useful low-dimensional summary of each storm event for extreme value purposes.

Understanding the interaction of ocean environments with fixed and floating marine structures is critical to the design of offshore and coastal facilities. Structural response to environmental loading is typically the combined effect of multiple environmental variables over a period of time, and is the subject of extensive academic study. Knowledge of the tails of marginal and joint distributions of storm peak significant wave height and associated peak period is central to estimation, by simulation or closed form approximation, of the statistics of extreme structural response, and hence of structural reliability and safety. The directional variability of extreme environments is also important to estimate, particularly for complex responses of floating structures, but also for the design of safety-critical top-side facilities on fixed structures (see, for example, Tromans and Vanderschuren 1995, Winterstein et al. 1993 and Ewans and Jonathan 2014).

The layout of the paper is as follows. In Section 2, we introduce the motivating oceanographic application from the northern North Sea. Section 3 is a description of marginal inference (incorporating threshold and quantile estimation using non-crossing quantile regression, generalised Pareto modelling of threshold exceedances and transformation to the Gumbel scale), and conditional extremes inference incorporating spline covariate effects. Section 4 evaluates the model in application to simulated samples with known

covariate-dependent extremal characteristics. Section 5 addresses estimation of conditional extremes models for the motivating application under consideration, including full bootstrap uncertainty analysis. A brief summary of a second application to the South Atlantic Ocean is also provided. Section 6 discusses further enhancement to include multidimensional covariate effects, and illustrates the dependence structure of conditional extremes model parameter estimates for the northern North Sea application. Motivation and algorithmic details for estimation of non-crossing quantiles are given in the appendix (Section 7).

2. MOTIVATING APPLICATION

We motivate and illustrate the methodology by considering joint estimation of extreme values of storm peak H_S and T_P with directional covariate at a location in the northern North Sea. The sample for analysis corresponds to hindcast values of storm peak H_S over threshold, observed during periods of storm events, and associated values for T_P , together with corresponding storm peak directions (strictly, the dominant wave direction at storm peak). All references to H_S and T_P below are to storm peak events.

The wind and wave climate of the North Sea is highly variable both temporarily and spatially. Main climatic variation is from north to south. Low pressure systems travelling from west to east across the Atlantic cross the northern North Sea, resulting in more intense winds and waves. Wind and wave directions rotate, however, through most compass headings during the passage of storms. In the southern North Sea, these storms produce winds mainly from the northwest to northeast and waves largely from the north. The northern North Sea is also more exposed to long period swell waves from the North Atlantic and Norwegian Sea. Some swell energy may also propagate into the southern North Sea. Throughout most of the North Sea, the land masses of the UK, Scandinavia, and Western Europe introduce fetch-limited effects on wave generation, resulting in less extreme conditions for wind and wave

directions corresponding to upwind land shadows. For this reason also, the ocean environment in the southern North Sea is relatively less severe.

The northern North Sea sample consists of 1163 pairs of values for the period October 1964 to August 1998 . With direction from which a storm travels expressed in degrees clockwise with respect to north, Figure 1 consists of scatter plots of H_S (upper panel) and T_P (lower panel) versus direction. The influence of relatively long fetches corresponding to the Atlantic Ocean (directional sector [230, 280)), the Norwegian Sea (directional sector [320, 20)) and the North Sea (directional sector [140, 200)) are clear, as is that of the land shadow of Norway (directional sector [20, 140)). Scatter plots of H_S (horizontal) and T_P (vertical) for the directional sectors identified in Figure 1 are given in Figure 2. The $H_S - T_P$ dependence for Atlantic storms is different to that for storms emanating from the south (i.e. travelling up the North Sea), suggesting that the dependence between H_S and T_P varies with storm direction. Inspection of samples from other ocean basins (see, for example, Section 5.3) suggests that the dependence between H_S and T_P will be, in general, a function of storm direction.

[Figure 1 about here.]

[Figure 2 about here.]

3. MODEL

Consider two random variables $\dot{X}_1(\theta), \dot{X}_2(\theta)$ of common covariate θ . We seek to characterise their joint extremal structure for any particular value of the covariate θ . We assume that the joint tail of $\dot{X}_1(\theta)$ and $\dot{X}_2(\theta)$ can be characterised adequately using covariate θ . The framework developed here assumes that θ is multidimensional. For example, in an oceanographic context, we might be interested in a spatio-directional description of the joint extremal behaviour of H_S and associated T_P , in which case we might adopt latitude, longitude

and storm direction as a three-dimensional covariate. Simulation studies and applications below assume a single covariate for ease of description and illustration.

We now describe the extension of the conditional extremes modelling procedure, outlined in Section 1, to incorporate covariates. Sections 3.1 and 3.2 outline the coupled marginal generalised Pareto modelling of threshold exceedances and quantile regression for threshold selection respectively. Section 3.3 discusses transformation to standard Gumbel scale, necessary for application of the conditional extremes model in Section 3.4. The adoption of periodic B-spline representations for model parameter functions is outlined in Section 3.5. Section 3.6 describes the simulation procedure used to generate samples of conditional extremes of one variate given large values of the other. The bootstrap procedure used to estimate model parameter and return value uncertainty is outlined in Section 3.7.

3.1. Generalised Pareto model for threshold exceedances

Marginally, for each of $\dot{X}_1(\theta)$, $\dot{X}_2(\theta)$ in turn for a given value of θ , we assume that, conditional on exceeding a large value, the corresponding random variables are generalised Pareto distributed:

$$\Pr(\dot{X}_j(\theta) > \dot{x} | \dot{X}_j(\theta) > \psi_j(\theta; \tau_{j*})) = \left(1 + \frac{\xi_j(\theta)}{\zeta_j(\theta)} (\dot{x} - \psi_j(\theta; \tau_{j*}))\right)^{-1/\xi_j(\theta)} \text{ for } j = 1, 2$$

for $\dot{x} > \psi_j(\theta; \tau_{j*})$, $(1 + \frac{\xi_j(\theta)}{\zeta_j(\theta)} (\dot{x} - \psi_j(\theta; \tau_{j*}))) > 0$ and $\zeta_j(\theta) > 0$. $\psi_j(\theta; \tau_{j*})$ is a pre-selected quantile threshold, assumed to be a smooth function of θ , associated with a non exceedance probability τ_{j*} :

$$\Pr(\dot{X}_j(\theta) \leq \psi_j(\theta; \tau_{j*})) = \tau_{j*}$$

Model parameters $\xi_j(\theta)$, $\zeta_j(\theta)$, respectively generalised Pareto shape and scale, are also assumed to be smooth functions of the covariate.

For a sample of values $\{x_{ij}\}_{i=1}^n$, $j = 1, 2$, corresponding to set $\{\theta_i\}_{i=1}^n$ of known covariate values, and pre-specified threshold $\psi_j(\theta; \tau_{j*})$, estimates for the values of the functions $\xi_j(\theta)$ and $\zeta_j(\theta)$ at $\{\theta_i\}_{i=1}^n$ can be obtained in principle by maximum likelihood estimation, by minimising the negative log-likelihood:

$$\ell_{GP,j} = \sum_{i=1}^n \log \zeta_j(\theta_i) + \frac{1}{\xi_j(\theta_i)} \log \left(1 + \frac{\xi_j(\theta_i)}{\zeta_j(\theta_i)} (x_{ij} - \psi_j(\theta_i; \tau_{j*})) \right) \text{ for } j = 1, 2$$

Each of the parameter functions $\psi_j(\theta; \tau_{j*})$, $\xi_j(\theta)$ and $\zeta_j(\theta)$ can be specified as a linear combination of suitable basis functions, such as periodic splines and Fourier series for periodic covariates such as direction, as discussed in Section 3.5. In this case, we regulate parameter smoothness with covariate using a penalised likelihood fitting criterion:

$$\ell_{GP,j}^* = \ell_{GP,j} + \lambda_{\xi_j} R_{\xi_j} + \lambda_{\zeta_j} R_{\zeta_j} \text{ for } j = 1, 2$$

for roughness coefficients λ_{ξ_j} , λ_{ζ_j} , and parameter roughnesses R_{ξ_j} , R_{ζ_j} which are easily evaluated for suitable choice of basis. The values of roughness coefficients are selected using cross-validation to maximise the predictive performance of the model, measured in terms of the sum, over cross-validation iterations, of the likelihood value for the omitted data at each iteration.

3.2. Quantile regression model for thresholds

For each random variable in turn, the quantile threshold $\psi_j(\theta; \tau)$ corresponding to quantile probability τ is estimated using quantile regression, by minimising the roughness penalised loss criterion:

$$\ell_{QR,j}^* = \left\{ \tau \sum_{i, r_{ij} \geq 0}^n |r_{ij}| + (1 - \tau) \sum_{i, r_{ij} < 0}^n |r_{ij}| \right\} + \lambda_{\psi_j} R_{\psi_j} \text{ for } j = 1, 2$$

and residuals $r_{ij} = \dot{x}_{ij} - \psi_j(\theta_i; \tau)$. The terms in brackets on the right hand side correspond to the unpenalised quantile regression loss criterion. Parameter roughness R_{ψ_j} can be evaluated in closed form for efficient estimation. The value of roughness coefficient λ_{ψ_j} is selected using cross-validation to maximise the predictive performance of the quantile regression model.

In practice, quantile regression thresholds are estimated for an increasing sequence of D quantile probabilities $0 < \tau_1 < \tau_2 < \dots < \tau_d < \dots < \tau_D < 1$. For each choice of τ_d , standard diagnostic plots for generalised Pareto fitting (such as the variation of the estimated shape parameter or some extreme quantile estimate with threshold) are examined. The lowest value of quantile probability consistent with an adequate generalised Pareto fit is selected as τ_{j*} . The quantile regression threshold estimates for quantile probabilities $\leq \tau_{j*}$ are useful for marginal transformation to standard Gumbel scale, discussed in the next section.

Using the approach described in Bollaerts (2009) and Bollaerts et al. (2006) developed from the work of Marx and Eilers, and Koenker (see, for example, Marx and Eilers 1998, Koenker 2005, Eilers and Marx 2010), quantile regression can be formulated as a linear programme with roughness penalisation for simultaneous estimation of all quantile levels with the non-crossing constraint that $\psi_j(\theta; \tau_d) \leq \psi_j(\theta; \tau_{d+1})$, for $d = 1, 2, \dots, D - 1, j=1,2$. Motivation for and details of the linear programme are given in Section 7.1.

3.3. Marginal transformation to standard Gumbel scale

The conditional extremes model is applied to random variables with standard Gumbel marginal distributions. For each random variable in turn, the quantile regression models for different quantile probabilities, and the marginal generalised Pareto model for threshold exceedances, provide a means to transform from sample $\{\dot{x}_{ij}\}_{i=1}^n$ corresponding to random variable $\dot{X}_j(\theta)$ at the set $\{\theta_i\}_{i=1}^n$ of known covariate values to an equivalent sample $\{x_{ij}\}_{i=1}^n$ corresponding to random variable $X_j(\theta)$ with standard Gumbel distribution for any θ .

Above the threshold $\psi_j(\theta; \tau_{j^*})$, the unconditional cumulative distribution function for threshold exceedances $\dot{x} > \psi_j(\theta; \tau_{j^*})$, for any value of θ , is given by:

$$\Pr(\dot{X}_j(\theta) \leq \dot{x}) = 1 - (1 - \tau_{j^*}) \Pr(\dot{X}_j(\theta) > \dot{x} | \dot{X}_j(\theta) > \psi_j(\theta; \tau_{j^*})) \text{ for } j = 1, 2$$

Below the threshold, in the absence of a parametric form for the cumulative distribution function, we approximate it using:

$$\Pr(\dot{X}_j(\theta) \leq \dot{x}) \approx \tau_d + (\tau_d - \tau_{d-1}) \frac{(\dot{x} - \psi_j(\theta; \tau_{d-1}))}{(\psi_j(\theta; \tau_d) - \psi_j(\theta; \tau_{d-1}))} \text{ for } j = 1, 2$$

where $\psi_j(\theta; \tau_{d-1}) \leq \dot{x} < \psi_j(\theta; \tau_d)$ for the sequence of quantile probabilities τ_d such that $\tau_d \leq \tau_{j^*}$.

Using the probability integral transform, we can transform from original to standard Gumbel scales, since:

$$\Pr(X_j(\theta) \leq x) = \exp(-\exp(-x)) = \Pr(\dot{X}_j(\theta) \leq \dot{x}) \text{ for } j = 1, 2$$

so that the individuals in the transformed sample $\{x_{ij}\}_{i=1}^n, j = 1, 2$ are given by:

$$x_{ij} = -\log(-\log(\Pr(X_j(\theta_i) \leq x))) \text{ for } i = 1, 2, \dots, n, \text{ and } j = 1, 2$$

now assumed to be marginally stationary.

3.4. Conditional extremes model

For positively dependent random variables $X_1(\theta), X_2(\theta)$ with standard Gumbel marginal distributions for any θ , we extend the asymptotic argument of Heffernan and Tawn (2004) for the form of the conditional distribution of $X_{j^c}(\theta), j^c = 1, 2$ given the value of $X_j(\theta), j = 1, 2, j \neq j^c$ for any value of covariate θ :

$$(X_{j^c}(\theta) | X_j(\theta) = x) = \alpha_j(\theta)x + x^{\beta_j(\theta)}W_j(\theta) \text{ for } j, j^c = 1, 2, j^c \neq j$$

and $x > \phi_j(\kappa_{j*})$ where $\phi_j(\kappa_{j*})$ is a threshold with non-exceedance probability κ_{j*} above which the conditional extremes model fits well. The parameter functions $\alpha_j(\theta) \in [0, 1]$, $\beta_j(\theta) \in (-\infty, 1]$ vary smoothly with covariate θ . $W_j(\theta)$ is a random variable drawn from an unknown distribution. We assume that the standardised variable $Z_j = (W_j(\theta) - \mu_j(\theta))/\sigma_j(\theta)$ follows a common distribution G_j , independent of covariate, for smooth location and scale parameter functions $\mu_j(\theta)$, $\sigma_j(\theta) > 0$. We write, for any value of θ :

$$(X_{j^c}(\theta)|X_j(\theta) = x) = \alpha_j(\theta)x + x^{\beta_j(\theta)}(\mu_j(\theta) + \sigma_j(\theta)Z_j) \text{ for } j, j^c = 1, 2, j^c \neq j$$

For potentially negatively dependent variables (corresponding to $\alpha = 0$ and $\beta < 0$, see Heffernan and Tawn 2004), extended forms of the equations above are available in the stationary case, and can be easily specified for the non-stationary case also. However, for most applications of practical interest, and for oceanographic applications involving H_S and T_P in particular, we expect $\alpha > 0$ so that positive dependence only need be considered. In practice, were we to estimate α to be approximately zero, the extended form of the conditional extremes model would then need to be considered. To estimate the parameter functions $\alpha_j(\theta)$, $\beta_j(\theta)$, $\mu_j(\theta)$ and $\sigma_j(\theta)$, we follow Heffernan and Tawn (2004) in assuming that G_j is the standard normal distribution. The corresponding negative log likelihood for pairs $\{x_{i1}, x_{i2}\}$ from the original sample for which $x_{ij} > \phi_j(\kappa_{j*})$, conditioned on $X_j(\theta)$ is:

$$\ell_{CE,j} = \sum_{i, x_{ij} > \phi_j(\theta_i; \kappa_{j*})} \log s_{ij} + \frac{(x_{ij^c} - m_{ij})^2}{2s_{ij}^2} \text{ for } j, j^c = 1, 2, j^c \neq j$$

where $m_{ij} = \alpha_j(\theta_i)x_{ij} + \mu_j(\theta_i)x_{ij}^{\beta_j(\theta_i)}$ and $s_{ij} = \sigma_{ij}x_{ij}^{\beta_j(\theta_i)}$. Adopting a penalisation procedure to regulate parameter roughness, the penalised negative log likelihood is:

$$\ell_{CE,j}^* = \ell_{CE,j} + \lambda_{\alpha_j}R_{\alpha_j} + \lambda_{\beta_j}R_{\beta_j} + \lambda_{\mu_j}R_{\mu_j} + \lambda_{\sigma_j}R_{\sigma_j} \text{ for } j = 1, 2$$

where parameter roughnesses R_{α_j} , R_{β_j} , R_{μ_j} , R_{σ_j} are easily evaluated, and roughness coefficients λ_{α_j} , λ_{β_j} , λ_{μ_j} , λ_{σ_j} are estimated using cross-validation. ‘To reduce computational burden, we choose to fix the relative size of the roughness coefficients so that only one roughness coefficient λ_j ($= \delta_{\alpha_j}\lambda_{\alpha_j} + \delta_{\beta_j}\lambda_{\beta_j} + \delta_{\mu_j}\lambda_{\mu_j} + \delta_{\sigma_j}\lambda_{\sigma_j}$) is estimated for each $j = 1, 2$. The values of δ_{α_j} , δ_{β_j} , δ_{μ_j} and δ_{σ_j} used are set by careful experimentation, as discussed in Section 4.1. Residuals:

$$r_{ij} = \frac{1}{\widehat{\sigma}_j(\theta_i)} ((x_{ij^c} - \widehat{\alpha}_j(\theta_i)x_{ij})x_{ij}^{-\widehat{\beta}_j(\theta_i)} - \widehat{\mu}_j(\theta_i)) \text{ for } j, j^c = 1, 2, j^c \neq j$$

evaluated for $x_{ij} > \phi_j(\kappa_{j*})$ are inspected to confirm reasonable model fit, as discussed in Section 4. The set of residuals is also used as a random sample of values for Z_j from the unknown distribution G_j for simulation to estimate extremes quantiles in Section 5.

The limit assumption underlying the non-stationary conditional extremes model is itself an extension of that made by Heffernan and Tawn (2004), and has been reported previously (Jonathan et al. 2013). For positively dependent random variables $X_1(\theta)$ and $X_2(\theta)$ of common covariate θ with standard Gumbel marginal distributions for any value of θ , we assume that the standardised variable:

$$Z_j = \sigma_j(\theta)^{-1} \left(\frac{X_{j^c}(\theta) - \alpha_j(\theta)}{x_j^{\beta_j(\theta)}} - \mu_j(\theta) \right) \text{ for } j, j^c = 1, 2, j^c \neq j$$

is such that:

$$\Pr(Z_j \leq z | X_j = x_j, \theta) \rightarrow G(z) \text{ as } z \rightarrow \infty$$

for some non-degenerate distribution G independent of covariate θ , where α , β , μ and σ are smooth functions of θ . For the current simulation studies and applications, this would seem to be a reasonable assumption, by construction for simulations, and from physical considerations and inspection of model diagnostics for applications.

3.5. Parameter functional forms

Physical considerations usually suggest that parameters $\psi_j(\theta)$, $\xi_j(\theta)$, $\zeta_j(\theta)$, $\alpha_j(\theta)$, $\beta_j(\theta)$, $\mu_j(\theta)$ and $\sigma_j(\theta)$ would be expected to vary smoothly with respect to covariates θ . In some situations, for example involving directional or seasonal covariates, we might also expect parameters to vary periodically with respect to some components of θ . Adopting the notation $\eta(\theta)$ for a typical parameter function, this can be achieved by expressing $\eta(\theta)$ in terms of an appropriate basis for the covariate domain. In the illustrations reported here, with periodic covariate on $[0, 360)$, we adopt a basis of periodic B-splines of appropriate order (cubic throughout this work). We calculate the B-spline basis matrix B ($m \times p$) for an index set of m (typically less than sample size, n) covariate values, at p uniformly spaced knot locations on $[0, 360)$. More generally, for example in the case of a spatio-directional covariate, we would define B-spline bases B_{Lng} ($m_{Lng} \times p_{Lng}$), B_{Ltt} ($m_{Ltt} \times p_{Ltt}$) and B_{Drc} ($m_{Drc} \times p_{Drc}$) for longitude, latitude and direction respectively on the domain of interest, from which the full spatio-directional basis B ($m \times p = m_{Lng}m_{Ltt}m_{Drc} \times p_{Lng}p_{Ltt}p_{Drc}$) is evaluated as

$$B = B_{Drc} \otimes B_{Ltt} \otimes B_{Lng}$$

where \otimes represents the Kronecker product. Values of η on the index set can then be expressed as

$$\eta = B\beta$$

for some vector β ($p \times 1$) of basis coefficients to be estimated.

The roughness R of $\eta(\theta)$ is evaluated on the index set using the approach of Eilers and Marx (2010). Writing the vector of differences of consecutive values of β as $\Delta\beta$, and vectors of second and higher order differences using $\Delta^k\beta = \Delta(\Delta^{k-1}\beta)$, $k > 1$, the roughness R of β

is given by

$$R = \beta' P \beta$$

where $P = (\Delta^k)'(\Delta^k)$ for differencing at order k . We use $k = 1$ throughout this work.

3.6. Simulation procedure

Inferences (for example, concerning the probabilities of extreme sets in $X_1(\theta)$ and $X_2(\theta)$, for some value or interval of θ) are drawn using the non-stationary conditional extremes approach by simulating joint occurrences using the estimated marginal and conditional extremes models. To simulate a realisation of an exceedance of a high quantile of the conditioning variate $X_j(\theta)$ ($j = 1, 2$) and a corresponding value of the conditioned variate $X_{j^c}(\theta)$ ($j^c = 1, 2, j^c \neq j$), for some θ , we proceed as follows:

1. Draw a value of covariate θ_s (for example, at random from the original sample or from some estimate of its distribution estimated from the sample),
2. Draw a value of residual r_{sj} from the set of residuals obtained during model fitting (see Section 3.4),
3. Draw a value x_{sj} of the conditioning variate from its standard Gumbel distribution,
4. If the value x_{sj} exceeds $\phi_j(\kappa_{j^*})$ continue to step 5, else repeat steps 1 to 3,
5. Estimate the value of the conditioned variate x_{sj^c} using:

$$x_{sj^c} = \hat{\alpha}_j(\theta_s)x_{sj} + x_{sj}^{\hat{\beta}_j(\theta_s)}(\hat{\mu}_j(\theta_s) + \hat{\sigma}_j(\theta_s)r_{sj})$$

where the obvious notation is used for the estimated values of model parameters (see Section 3.4),

6. Transform the pair x_{sj}, x_{sj^c} in turn to the original scale (to $\dot{x}_{sj}, \dot{x}_{sj^c}$) using the probability integral transform (see Section 3.3)

Note that this procedure can be extended to include realisations for which $x_{sj} \leq \phi_j(\kappa_{j*})$, by drawing a pair of values (for the conditioning and conditioned variates) at random from the subset of the Gumbel-transformed original sample (for which the conditioning variate is $\leq \phi_j(\kappa_{j*})$) at step 4.

3.7. Bootstrap estimation of parameter uncertainty

Throughout this work, 95% bootstrap confidence intervals for model parameters and return values are estimated using the bias-corrected and accelerated (BC_a) method (Efron 1987, DiCiccio and Efron 1996) by repeating the full analysis for 2000 bootstrap resamples of the original sample (for the simulation studies in Section 4) and the original storm peak data (for the oceanographic application in Section 5). In particular, estimation of optimal roughness penalties is performed independently for each bootstrap resample, so that uncertainty bands also reflect variability in these choices. It was also confirmed that 2000 resamples was sufficient to ensure stability of bootstrap confidence intervals.

4. SIMULATION STUDY

We evaluate the method in estimating extreme quantiles using samples drawn from known bivariate distributions with particular extremal dependence structures varying as a function of single periodic covariate $\theta \in [0, 360)$. Standard Gumbel marginal distributions are assumed throughout, so that marginal estimation of generalised Pareto parameters (Section 3.3) is unnecessary. Sample size is set to 6000, and modelling threshold set at the 90th percentile of the standard Gumbel distribution, so that a sample of size 600 is available for modelling, typical of samples of ocean extremes.

4.1. Case 1: Bivariate distribution with Normal dependence transformed marginally to standard Gumbel

For any value of covariate θ ,

$$(X_1(\theta), X_2(\theta)) = -\log(-\log(\Phi_{\Sigma(\theta)}(X_{1N}, X_{2N})))$$

where (X_{1N}, X_{2N}) are Normally-distributed random variables with zero mean and covariance matrix $\Sigma(\theta)$. $\Sigma(\theta)$ has elements $\Sigma_{11} = \Sigma_{22} = 1$, and $\Sigma_{12} = \Sigma_{21} = \rho(\theta)$. $\Phi_{\Sigma(\theta)}$ is the cumulative distribution function of $N(0, \Sigma(\theta))$. Conditionally, that $(X_{j^cN}(\theta)|X_{jN}(\theta) = x) \sim N(\rho(\theta)x, (1 - \rho^2(\theta)))$, for $j, j^c = 1, 2, j^c \neq j$ is useful for simulation of realisations. X_1 and X_2 have standard Gumbel marginal distributions, and are dependent (unless $\rho=0$) but asymptotically independent (unless $\rho=1$). Conditionally, $(X_{j^c}(\theta)|X_j(\theta) = x) = \rho^2(\theta)x + x^{1/2}W_j(\theta)$ for large x , for $j, j^c = 1, 2, j^c \neq j$, and W_j is (asymptotically) Normally-distributed (see, for example, Heffernan and Tawn 2004). The true values of $\alpha_j(\theta)$ and $\beta_j(\theta)$ in the conditional model (see Section 3.4) are therefore $\rho^2(\theta)$ and 0.5 respectively. In each of 6 intervals $[60 * (j - 1), 60 * j)$, $j = 1, 2, 3, \dots, 6$, of covariate θ , we set the value of $\rho^2(\theta)$ to be respectively 0.6, 0.9, 0.5, 0.1, 0.7 and 0.3 . We sample 1000 values at random from each interval to construct the sample. When aggregated across direction, a scatter plot (not shown) shows no evidence of heterogeneity. Yet when we partition the data with respect to covariate, as illustrated in Figure 3, the changing dependence structure of the sample is obvious.

[Figure 3 about here.]

The conditional extremes model is estimated assuming a cubic B-spline basis with ($p =$) 60 knots on $[0, 360)$, with a first order spline penalty ($k = 1$), as outlined in Section 3. 10-fold cross-validation is used to estimate the value of the single parameter roughness λ (see Section 3.4, suppressing the j subscript without loss of generality due to symmetry of X_1

and X_2), with values of weights $\delta_\alpha, \delta_\beta, \delta_\mu$ and δ_σ set to 1,1,10,10 respectively (following some experimentation). Resulting parameter estimates were found to be relatively insensitive to reasonable choices of weights.

Selection of optimal λ is illustrated in Figure 4. The resubstitution curve illustrates that increasing parameter flexibility with respect to covariate, by reducing λ , improves the model's ability to describe the sample. However, the cross-validation curve indicates a clear optimal choice for λ (of approximately 2.2 on the \log_{10} scale) to maximise the predictive performance of the model.

[Figure 4 about here.]

Figure 5 shows 2.5%, median and 97.5% bootstrap estimates for parameters α, β, μ and σ , together with the known theoretical values for α and β . There is good agreement in general, particularly in light of the step changes in the values of the true parameters with covariate.

[Figure 5 about here.]

Figure 6 shows bootstrap estimates for conditional return values of the associated variate X_2 , corresponding to values of the conditioning variate X_1 with non-exceedance probabilities (for period of sample) of (a) 0.99 and (b) 0.999, together with sample estimate and known value. Again, there is good agreement.

[Figure 6 about here.]

4.2. Case 2: Mixture of bivariate distribution with Normal dependence transformed marginally to standard Gumbel, and bivariate extreme value distribution with exchangeable logistic dependency and Gumbel marginal distributions

For this case, the functional form of the dependence between variables $X_1(\theta)$ and $X_2(\theta)$ changes as a function of θ , over the same six covariate sectors as defined for simulation Case

1. In the first, second and third sectors, with $\theta \in [0, 180)$, the dependence structure of case 1 is retained, with $\rho^2(\theta)$ set to 0.8, 0.1 and 0.8 respectively. In the remaining three sectors, $\theta \in [180, 360)$, a logistic dependence structure is assumed, such that in these sectors

$$\Pr(X_1(\theta) \leq x_1, X_2(\theta) \leq x_2) = \exp(-(\exp(x/\omega(\theta)) + \exp(y/\omega(\theta)))^{\omega(\theta)}),$$

with parameter $\omega(\theta)$ set to 0.6, 0.1 and 0.6 respectively. Scatter plots of the sample drawn per covariate sector are shown in Figure 7.

[Figure 7 about here.]

$X_1(\theta)$ and $X_2(\theta)$ have standard Gumbel marginal distributions for any choice of $\theta \in [0, 360)$. Conditionally, for $\theta \in [180, 360)$, $\Pr(X_2(\theta) \leq x_2 | X_1(\theta) = x_1)$ can be written in closed form. Sample realisations from this bivariate distribution are obtained by first sampling one variate from its marginal distribution and subsequently sampling the second variate from the conditional distribution given the value of the first variate. $X_1(\theta)$ and $X_2(\theta)$ are dependent (unless $\omega(\theta) = 1$). Conditionally, $(X_2(\theta) | X_1(\theta) = x_1) = x_1 + W(\theta)$ for sufficiently large x_1 , exhibiting asymptotic dependence (since $X_2(\theta) | X_1(\theta) = x_1$ depends on x_1 for large x_1). For $\theta \in [180, 360)$, the value of $\omega(\theta)$ ($\omega(\theta) < 1$) has no effect on the asymptotic conditional dependence structure and does not influence the parameters $\alpha(\theta)$ and $\beta(\theta)$ which are 0 and 1 respectively. The conditional extremes model was estimated using the procedure described for simulation Case 1. Figure 8 shows 2.5%, median and 97.5% bootstrap estimates for parameters α, β, μ and σ , together with the known theoretical values for α and β . Estimation of α is good. It is however noteworthy that identification of β is less precise for $\theta \in [180, 360)$, due to redundancy between β, μ and σ . Nevertheless, bootstrap estimates (in Figure 9) for conditional return values of the associated variate X_2 , corresponding to values of the conditioning variate X_1 with non-exceedance probabilities (for period of sample) of 0.99 and 0.999, together with sample estimate and known value, show good agreement.

[Figure 8 about here.]

[Figure 9 about here.]

5. APPLICATION

In this section we outline the application of the model introduced in Section 3 to the sample of northern North Sea storm peak significant wave heights and associated peak periods described in Section 2, with storm direction as the covariate of interest.

5.1. Transformed covariate scale

Marginal quantile regression thresholds $\psi_j(\theta; \tau_d)$ for quantile probabilities τ_d of 0.1, 0.2, ..., 0.9 were estimated for storm peak H_S ($j = 1$) and T_P ($j = 2$), and used in turn for generalised Pareto modelling. We found it advantageous to use evenly-spread transformed covariate values $\{\theta_i^*\}_{i=1}^n$, with:

$$\theta_i^* = \frac{360}{n}(r(\theta_i) - 1) \text{ for } i = 1, 2, \dots, n$$

where $r(\theta_i)$ is the rank of θ_i in the set of covariates, namely the position of θ_i in the set of covariate values sorted in ascending order. The set $\{\theta_i^*\}_{i=1}^n$ is uniformly distributed on $[0, 360)$ by design, stabilising estimation on the transformed θ^* scale. Interpreted on the original θ scale, the transformation imposes greater smoothness in directional sectors which are less frequently observed, and allows greater threshold flexibility in more frequently observed sectors, in a natural way according to the rate of occurrence of events from different directions. Estimates on the transformed scale necessarily also provide a relatively poor sample description for intervals of the original covariate θ which are under-represented in the sample; that is, for $\theta \in [20, 140)$ for the northern North Sea application (see Figure 1). In this interval, spline knot spacing on the θ scale is relatively very coarse. Parameter estimates are therefore essentially interpolated between those for the nearest spline knot locations, and

are not informative for the intervals themselves in general. For this reason, we choose not to include parameter estimates for this sparsely-populated covariate interval.

5.2. Northern North Sea

We find that for thresholds corresponding to $\tau_d = 0.7$ and above, we have reasonable generalised Pareto parameter stability for both H_S and T_P . We therefore set $\tau_{1*} = \tau_{2*} = 0.7$ and adopt $\psi(\theta; \tau_{j*}), j = 1, 2$ as extreme value thresholds for H_S and T_P respectively. Plots of non-crossing quantile regression threshold estimates for τ_d of 0.1, 0.2, ..., 0.7 are given in Figure 10 for H_S and T_P respectively.

[Figure 10 about here.]

Figure 11 shows bootstrap 2.5%, 50% and 97.5% estimates for non-stationary extreme value threshold, generalised Pareto shape and scale for both H_S and T_P (in black) with storm direction. For comparison, the figure also shows the corresponding stationary estimates (in grey). The discrepancy is clear, particularly for storms from the south. The relative profiles of generalised Pareto shape and scale estimates with direction illustrate their characteristic negative dependence (see, for example, Scarrott and MacDonald 2012). The large reduction in the value of shape for H_S in [210,330) represents a shortening of the generalised Pareto tail with increasing angle. This is compensated to some extent by increasing scale, but not fully - the largest values of shape are positive, for severe Atlantic storms, indicating a generalised Pareto distribution unbounded to the right.

Scatter plots (not shown) of the transformed sample on Gumbel scale, partitioned by covariate, show differences in dependence structure similar to those shown on the original scale in Figure 2. Directional differences in the relationship between H_S and T_P persist, as can be seen, for example, from consideration of gradients of T_P on H_S in panels (b) and (e), which we expect to be reflected in the estimate of conditional extremes slope parameter α . Using a format similar to that of Figure 11, Figure 12 shows estimates for the non-stationary

conditional extreme model parameters with storm direction. Estimates of α are in the region of 0.5, and of β are negative. Differences between non-stationary (black) and stationary (grey) estimates are greatest for α , with higher values corresponding to storms from the Atlantic, the Norwegian Sea and in particular for storms emanating from the south. It is noteworthy that estimates for σ appear somewhat smaller for storms from the south than for Atlantic storms. This concurs with physical intuition, in that Atlantic storms are likely to contain swell components with longer peak periods, unlike storms from the south; this is evident in Figure 2. The general characteristics of parameter estimates in Figure 12 are similar to those reported in Jonathan et al. (2013), for a similar northern North Sea sample.

[Figure 11 about here.]

[Figure 12 about here.]

To examine whether the non-stationary model is a significant improvement on the stationary model, Figure 13 illustrates 95% bootstrap confidence intervals for the difference between non-stationary and stationary parameter estimates for the northern North Sea application. It can be seen that the confidence interval for the difference in parameter estimates for α does not include zero for storm direction of approximately 270° , indicating that the non-stationary model yields a significantly better representation for this parameter. For the other parameters, however, the corresponding confidence intervals do include zero.

[Figure 13 about here.]

Bootstrap estimates of marginal return values for T_P with storm direction, corresponding to non-exceedance probability of 0.999 relative to the period of the sample, are shown in Figure 14. The figure also gives bootstrap estimates for return values for T_P conditional on values of H_S with the same non-exceedance probabilities. For both marginal and conditional return values, corresponding return values assuming stationarity are shown for comparison (in grey). There is general agreement between stationary and non-stationary estimates for

both marginal and conditional return values, the latter reduced with respect to the former due to the dependence between T_P and H_S . There is evidence for direction trends in return values, particularly for Atlantic storms compared to others. Inspection of plots corresponding to Figure 14 for the difference between non-stationary and stationary estimates for the extreme quantile of T_P and T_P given H_S (not shown) indicates that significant differences between non-stationary and stationary estimates occur near 200° and 300° .

[Figure 14 about here.]

5.3. South Atlantic Ocean

A second application to a South Atlantic Ocean location, summarised briefly here, provides qualitatively similar findings to the northern North Sea application above. The sample of storm peak significant wave heights and associated storm peak periods, for the period June 1984 to July 1995, can be partitioned approximately by storm direction into three subsets by consideration of fetch conditions. Inspection of scatter plots for T_P against H_S for these subsets suggests that storm direction influences dependence structure. A non-stationary conditional extremes model is estimated as for the northern North Sea application, with non-stationary extreme value threshold, marginal and dependence model parameters. As for the northern North Sea, conditional extremes parameter α again shows greatest variability with storm direction, but values for α are generally lower than for the northern North Sea. Corresponding non-stationary marginal and conditional return values for T_P show stronger directional variation than their northern North Sea counterparts, with local peaks corresponding to northern and southern Atlantic storms.

6. DISCUSSION

In this article, we introduce an extension of the conditional extremes model of Heffernan and Tawn (2004) facilitating general non-stationary conditional extremes inference using spline representations of model parameters with respect to multidimensional covariates. We evaluate and illustrate the approach for bivariate extremes with respect to a single periodic covariate. Simulation studies show that the methodology provides good estimation of return values when sample dependence characteristics vary with covariate. Estimation of model parameter α is generally relatively good. Identification of β is sometimes difficult due to degeneracy of β, μ and σ , as seen for simulation Case 2 (in Section 4.2) in particular. Applications to hindcast ocean storm time-series of storm peak significant wave height and associated peak period show that storm directional variability is evident, and that directional extreme values models are necessary in general for realistic characterisation of marginal and dependence structure, and for reliable estimation of return values. Non-stationarity of conditional return values for peak period is more prominent for the South Atlantic Ocean application. Parameter estimates from the conditional extremes model exhibit interesting dependence. For example, Figure 15 gives convex hulls enclosing estimates for pairs of conditional extremes parameters for 2000 bootstrap re-samples within each of 6 directional sectors (see Figure 2) for the northern North Sea application. Some features are noteworthy, such as the negative dependence between β and μ , and β and σ , and the positive dependence between μ and σ for this application.

[Figure 15 about here.]

We adopt penalised B-splines as flexible bases to describe smooth variation of model parameters with respect to covariates. Extension to higher dimensions is limited only by computational (memory) constraints in general. Using sparse matrix operations and innovative algorithms such as those proposed by Eilers and Marx (2010), computational

efficiency is vastly improved compared with naive implementation. Adaptation to a random field representation is relatively straightforward, given the algorithmic similarity of the two approaches. The choice of splines, as opposed to a Fourier description of parameter variation with covariates has two advantages. Firstly, B-splines have local support facilitating efficient and stable inference. Secondly, in contrast to Fourier, spline models are easily and consistently extendible to multidimensional covariates using Kronecker products of spline bases as discussed in Section 3.

The methodology requires a means for transformation of non-stationary marginal distributions to standard stationary form. Non-crossing quantile regression is one means of achieving this, at least approximately. Since examination of parameter estimates with respect to non-stationary thresholds corresponding to different non-exceedance probabilities is a critical diagnostic in extreme value analysis generally, we recommend non-crossing quantile regression as a useful approach to threshold selection when covariate effects are thought to be influential. Reliable quantification of parameter and return value uncertainty is similarly important. In this work, we use an all-encompassing bootstrap scheme, re-sampling the original storm peak sample with replacement, repeating the full analysis including cross-validatory assessment of optimal parameter roughness for each re-sample, thereby capturing uncertainties throughout the modelling procedure.

ACKNOWLEDGEMENT

The authors acknowledge discussions with Graham Feld, Michael Vogel, Yanyun Wu and Elena Zanini at Shell, and Ross Towe of Lancaster University, UK. The authors further acknowledge thorough and thoughtful comments from anonymous reviewers and associate editor.

7. APPENDIX

7.1. Non-crossing quantile regression

Quantile regression estimation minimising the roughness penalised loss criterion from Section 3.2 can be achieved using linear programming. A linear programme for multiple simultaneous non-crossing quantile regression is outlined in this section, preceded by a brief motivation.

Suppressing the indexing subscript j for conciseness (so that $\dot{x}_i = \dot{x}_{ij}$, $i = 1, 2, \dots, n$, for the j of interest is understood) and temporarily ignoring roughness penalisation, a simplified loss criterion for estimation of quantile $\psi(\theta; \tau)$ with non-exceedance probability τ becomes

$$\ell_{QR}^* = \left\{ \tau \sum_{i, r_i \geq 0}^n |r_i| + (1 - \tau) \sum_{i, r_i < 0}^n |r_i| \right\}$$

with $r_i = \dot{x}_i - \psi(\theta_i; \tau)$. Minimisation of this expression can be written as

$$\min_{\psi} \left\{ \tau \sum_i^n u_i^+ + (1 - \tau) \sum_i^n u_i^- \right\}$$

subject to

$$\widehat{\psi}(\theta_i; \tau) + u_i^+ - u_i^- = \dot{x}_i \text{ for } i = 1, 2, \dots, n,$$

where u_i^+ and u_i^- are “positive and negative residuals” respectively defined as

$$\begin{aligned} u_i^+ &= |\psi(\theta_i; \tau) - \widehat{\psi}(\theta_i; \tau)| \text{ for } \psi(\theta_i; \tau) > \widehat{\psi}(\theta_i; \tau) \\ &= 0 \text{ otherwise, for } i = 1, 2, \dots, n, \end{aligned}$$

and

$$\begin{aligned} u_i^- &= |\psi(\theta_i; \tau) - \widehat{\psi}(\theta_i; \tau)| \text{ for } \psi(\theta_i; \tau) < \widehat{\psi}(\theta_i; \tau) \\ &= 0 \text{ otherwise, for } i = 1, 2, \dots, n. \end{aligned}$$

In this form, the quantile regression is easily estimated using linear programming. Reinstating roughness penalisation, the quantile regression can be expressed as

$$\min_{\psi} \left\{ \tau \sum_i^n u_i^+ + (1 - \tau) \sum_i^n u_i^- + \lambda_{\psi} \left(\sum_{j=k+1}^p (v_j^+ + v_j^-) \right) \right\}$$

subject to an additional constraint

$$\Delta^k \beta_j + v_j^+ - v_j^- = 0 \text{ for } j = k + 1, k + 2, \dots, p$$

where β is the p -vector of basis coefficients for the quantile function (such that, in matrix terms, $\psi = B\beta$, for spline basis matrix B), $\Delta^k \beta$ is its k th difference, defined for $j = k + 1, k + 2, \dots, p$, and v_j^+ and v_j^- are positive and negative residuals of the difference respectively, defined as

$$\begin{aligned} v_j^+ &= |\Delta^k \beta_j| \text{ for } \Delta^k \beta_j > 0 \\ &= 0 \text{ otherwise, for } j = k + 1, k + 2, \dots, p, \end{aligned}$$

and

$$\begin{aligned} v_j^- &= |\Delta^k \beta_j| \text{ for } \Delta^k \beta_j < 0 \\ &= 0 \text{ otherwise, for } j = k + 1, k + 2, \dots, p. \end{aligned}$$

In this form, roughness penalised quantile regression for a single quantile is also easily estimated using linear programming.

Estimation of non-crossing roughness penalised quantile regression for two quantile non-exceedance probabilities τ_1 and τ_2 ($> \tau_1$) can also be expressed as a linear programme, in which both quantile functions are estimated simultaneously subject to a further constraint

that

$$\beta_{\tau_2,j} - \beta_{\tau_1,j} + t_j = 0 \text{ for } j = k + 1, k + 2, \dots, p$$

with $t_j \geq 0$, for $j = k + 1, k + 2, \dots, p$, where β_{τ_1} and β_{τ_2} are p -vectors of basis coefficients.

In linear programming terms, simultaneous estimation of D non-crossing quantile functions can therefore be expressed as

$$\min_{\phi} c' \phi \text{ such that } A\phi = b$$

where

$$c = \left[Q_1; \dots Q_d; \dots Q_D; \begin{matrix} \mathbf{0} \\ (D-1)p \times 1 \end{matrix} \right],$$

$$Q_d = \left[\begin{matrix} \mathbf{0}; & \tau_d \mathbf{1}; & (1 - \tau_d) \mathbf{1}; & \lambda_\psi \mathbf{1}; & \lambda_\psi \mathbf{1} \\ p \times 1 & n \times 1 & n \times 1 & p \times 1 & p \times 1 \end{matrix} \right],$$

$$A = \left(A_1; A_2; A_3 \right)$$

and

$$b = \left(b_1; b_2; b_3 \right)$$

The linear combination $c' \phi$ is the loss criterion for simultaneous quantile regression, and pairs (A_1, b_1) , (A_2, b_2) and (A_3, b_3) in the constraint term respectively impose spline goodness of fit, smoothness and quantile non-crossing. The optimal value of the common roughness coefficient λ_ψ is estimated using cross-validation.

In detail,

$$\begin{aligned} \text{For spline fit: } A_1 &= \begin{bmatrix} \mathbf{I} & & & & \\ D \times D & \otimes & \begin{bmatrix} B & \mathbf{I} & -\mathbf{I} & \mathbf{0} & \mathbf{0} \\ n \times p & n \times n & n \times n & n \times (p-k) & n \times (p-k) \end{bmatrix} & \\ & & & & Dn \times (D-1)p \end{bmatrix}, \\ b_1 &= \begin{bmatrix} \dot{x}_1; & \dots & \dot{x}_d; & \dots & \dot{x}_D \\ n \times 1 & & n \times 1 & & n \times 1 \end{bmatrix}, \\ \dot{x}_d &= \{\dot{x}_i\}_{i=1}^n \text{ for } d = 1, 2, \dots, D, \end{aligned}$$

$$\begin{aligned} \text{For smoothness: } A_2 &= \begin{bmatrix} \mathbf{I} & & & & \\ D \times D & \otimes & \begin{bmatrix} \Delta^k & \mathbf{0} & \mathbf{0} & \mathbf{I} & -\mathbf{I} \\ (p-k) \times p & (p-k) \times n & (p-k) \times n & (p-k) \times (p-k) & (p-k) \times (p-k) \end{bmatrix} & \\ & & & & D(p-k) \times (D-1)p \end{bmatrix}, \\ b_2 &= \begin{bmatrix} \mathbf{0} \\ D(p-k) \times 1 \end{bmatrix}. \end{aligned}$$

$$\begin{aligned} \text{For non-crossing: } A_3 &= \begin{bmatrix} \Delta^1 & & & & \\ D-1 \times D & \otimes & \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ p \times p & p \times n & p \times n & p \times (p-k) & p \times (p-k) \end{bmatrix} & \\ & & & & (D-1)p \times (D-1)p \end{bmatrix}, \\ b_3 &= \begin{bmatrix} \mathbf{0} \\ (D-1)p \times 1 \end{bmatrix}. \end{aligned}$$

where the parameter vector ϕ is

$$\phi = \begin{bmatrix} \phi_1; & \dots & \phi_d; & \dots & \phi_D; & \mathbf{t} \\ & & & & & (D-1)p \times 1 \end{bmatrix}$$

and

$$\phi_d = \begin{bmatrix} \beta_{\psi_d}; & \mathbf{u}_d^+; & \mathbf{u}_d^-; & \mathbf{v}_d^+; & \mathbf{v}_d^- \\ p \times 1 & n \times 1 & n \times 1 & (p-k) \times 1 & (p-k) \times 1 \end{bmatrix}$$

where $\mathbf{u}_d^+ \geq 0$, $\mathbf{u}_d^- \geq 0$, $\mathbf{v}_d^+ \geq 0$, $\mathbf{v}_d^- \geq 0$ and $\mathbf{t} \geq 0$ are vectors of “residuals” (or *slack* variables in optimisation terminology) and β_{ψ_d} are spline coefficients for the d th quantile, $d = 1, 2, \dots, D$.

Illustrations of non-crossing quantile estimates for H_S and T_P from the northern North Sea application are given in Figure 10 of Section 5.2.

REFERENCES

- Anderson, C., Carter, D., Cotton, P., 2001. Wave climate variability and impact on offshore design extremes. Report commissioned from the University of Sheffield and Satellite Observing Systems for Shell International.
- Atyeo, J., Walshaw, D., 2009. A region-based hierarchical model for extreme rainfall over the UK, incorporating spatial dependence and temporal trend. *Environmetrics* 23, 509–521.
- Bollaerts, K., 2009. Statistical models in epidemiology and quantitative microbial risk assessment applied to salmonella in pork. Ph.D. thesis, Hasselt University, Belgium.
- Bollaerts, K., Eilers, P. H. C., Aerts, M., 2006. Quantile regression with monotonicity restrictions using P-splines and the L1 norm. *Statistical Modelling* 6, 189–207.
- Butler, A., Heffernan, J. E., Tawn, J. A., Flather, R. A., 2007. Trend estimation in extremes of synthetic North Sea surges. *J. Roy. Statist. Soc. C* 56, 395–414.
- Chavez-Demoulin, V., Davison, A., 2005. Generalized additive modelling of sample extremes. *J. Roy. Statist. Soc. Series C: Applied Statistics* 54, 207.
- Coles, S., Walshaw, D., 1994. Directional modelling of extreme wind speeds. *Applied Statistics* 43, 139–157.
- Coles, S. G., Casson, E., 1998. Extreme value modelling of hurricane wind speeds. *Structural Safety* 20, 283–296.
- Cooley, D., Naveau, P., Jomelli, V., Rabatel, A., Grancher, D., 2006. A Bayesian hierarchical extreme value model for lichenometry. *Environmetrics* 17, 555–574.
- Cox, A., Swail, V., 2000. A global wave hindcast over the period 1958-1997: Validation and climate assessment. *J. Geophys. Res. (Oceans)* 106, 2313–2329.
- Davison, A., Smith, R. L., 1990. Models for exceedances over high thresholds. *J. R. Statist. Soc. B* 52, 393.
- DiCiccio, T. J., Efron, B., 1996. Bootstrap confidence intervals. *Statist. Sci.* 11, 189–228.
- Eastoe, E. F., 2009. A hierarchical model for non-stationary multivariate extremes: a case study of surface-level ozone and NO_x data in the UK. *Environmetrics* 20, 428–444.
- Efron, B., 1987. Better bootstrap confidence intervals. *J. Am. Statist. Soc.* 82, 171–185.
- Eilers, P. H. C., Marx, B. D., 2010. Splines, knots and penalties. *Wiley Interscience Reviews: Computational Statistics* 2, 637–653.
- Ewans, K. C., Jonathan, P., 2008. The effect of directionality on Northern North Sea extreme wave design criteria. *J. Offshore Mechanics Arctic Engineering* 130, 10.

- Ewans, K. C., Jonathan, P., 2014. Evaluating environmental joint extremes for the offshore industry. *Journal of Marine Systems* 130, 124–130.
- Gilleland, E., Brown, B. G., Ammann, C. M., 2013. Spatial extreme value analysis to project extremes of large-scale indicators for severe weather. *Environmetrics* 24, 418–432.
- Grigg, O., Tawn, J., 2012. Threshold models for river flow extremes. *Environmetrics* 23, 295–305.
- Gyarmati-Szabo, J., Bogachev, L. V., Chen, H., 2011. Modelling threshold exceedances of air pollution concentrations via non-homogeneous Poisson process with multiple change-points. *Atmos. Environ.* 45, 5493–5503.
- Haver, S., Nyhus, K., 1986. A wave climate description for long term response calculations. *Proc. 5th OMAE Symp. IV*, 27–34.
- Heffernan, J. E., Resnick, S. I., 2007. Limit laws for random vectors with an extreme component. *Ann. Appl. Probab.* 17, 537–571.
- Heffernan, J. E., Tawn, J. A., 2004. A conditional approach for multivariate extreme values. *J. R. Statist. Soc. B* 66, 497.
- Jonathan, P., Ewans, K. C., 2011. Modelling the seasonality of extreme waves in the Gulf of Mexico. *ASME J. Offshore Mech. Arct. Eng.* 133:021104.
- Jonathan, P., Ewans, K. C., Flynn, J., 2012. Joint modelling of vertical profiles of large ocean currents. *Ocean Eng.* 42, 195–204.
- Jonathan, P., Ewans, K. C., Randell, D., 2013. Joint modelling of environmental parameters for extreme sea states incorporating covariate effects. *Coastal Engineering* 79, 22–31.
- Jonathan, P., Flynn, J., Ewans, K. C., 2010. Joint modelling of wave spectral parameters for extreme sea states. *Ocean Eng.* 37, 1070–1080.
- Keef, C., Papastathopoulos, I., Tawn, J. A., 2013a. Estimation of the conditional distribution of a vector variable given that one of its components is large: additional constraints for the Heffernan and Tawn model. *J. Mult. Anal.* 115, 396–404.
- Keef, C., Tawn, J., Svensson, C., 2009. Spatial risk assessment for extreme river flows. *J. Roy. Statist. Soc. C* 58, 601–618.
- Keef, C., Tawn, J. A., Lamb, R., 2013b. Estimating the probability of widespread flood events. *Environmetrics* 24, 13–21.
- Kinsman, B., 1984. *Wind waves: Their generation and propagation on the ocean surface*. Dover.

- Koenker, R., 2005. Quantile regression. Cambridge University Press.
- Kysely, J., Picek, J., Beranová, R., 2010. Estimating extremes in climate change simulations using the peaks-over-threshold method with a non-stationary threshold. *Global and Planetary Change* 72, 55–68.
- Marx, B. D., Eilers, P. H. C., 1998. Direct generalised additive modelling with penalised likelihood. *Computational statistics and data analysis* 28, 193–209.
- Mendez, F. J., Menendez, M., Luceno, A., Medina, R., Graham, N. E., 2008. Seasonality and duration in extreme value distributions of significant wave height. *Ocean Eng.* 35, 131–138.
- Northrop, P., Jonathan, P., 2011. Threshold modelling of spatially-dependent non-stationary extremes with application to hurricane-induced wave heights. *Environmetrics* 22, 799–809.
- Ruggiero, P., Komar, P. D., Allan, J. C., 2010. Increasing wave heights and extreme value projections: The wave climate of the US pacific northwest. *Coastal Eng.* 57, 539–522.
- Scarrott, C., MacDonald, A., 2012. A review of extreme value threshold estimation and uncertainty quantification. *REVSTAT - Statistical Journal* 10, 33–60.
- Stewart, R. H., 2008. Introduction to Physical Oceanography.
URL http://oceanworld.tamu.edu/resources/ocng_textbook/PDF_files/book.pdf
- Swail, V., Cox, A. T., 2000. On the use of NCEP/NCAR reanalysis surface marine wind fields for a long term North Atlantic wave hindcast. *J. Atmo. Tech.* 17, 532–545.
- Thompson, P., Cai, Y., Moyeed, R., Reeve, D., Stander, J., 2010. Bayesian nonparametric quantile regression using splines. *Computational Statistics and Data Analysis* 54, 1138–1150.
- Tromans, P. S., Vanderschuren, L., 1995. Risk based design conditions in the North Sea: Application of a new method. *Offshore Technology Conference, Houston (OTC-7683)*.
- Winterstein, S. R., Jha, A. K., Kumar, S., 1999. Reliability of floating structures: extreme response and load factor design. *Journal of Waterway, Port, Coastal, and Ocean Engineering* 125, 163–169.
- Winterstein, S. R., Ude, T. C., Cornell, C. A., Bjerager, P., Haver, S., 1993. Environmental parameters for extreme response: Inverse FORM with omission factors. In: *Proc. 6th Int. Conf. on Structural Safety and Reliability*, Innsbruck, Austria.

FIGURES




Figure 1. Northern North Sea sample. Directional scatter plots of (a) storm peak significant wave height and (b) associated storm peak period. Vertical grey lines indicate boundaries for 6 directional sectors, corresponding to different fetch conditions, estimated by inspection.

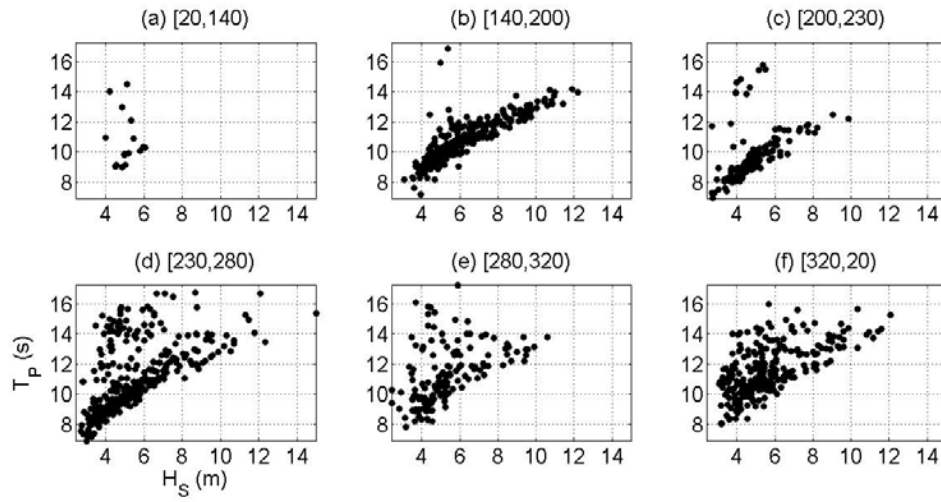


Figure 2. Northern North Sea sample. Scatter plots of associated storm peak period against storm peak significant wave height for the 6 directional sectors identified in Figure 1 (ordered, clockwise from 20°). The characteristics of dependence between T_P and H_S varies from sector to sector.

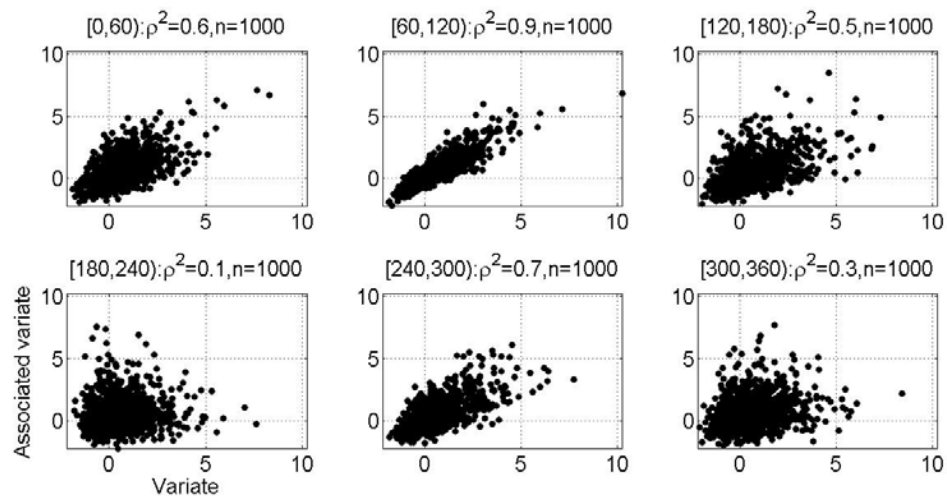


Figure 3. Simulation Case 1. Scatter plot per covariate interval. Values for intervals for covariate θ , parameter ρ^2 and sample size n are shown in each panel.

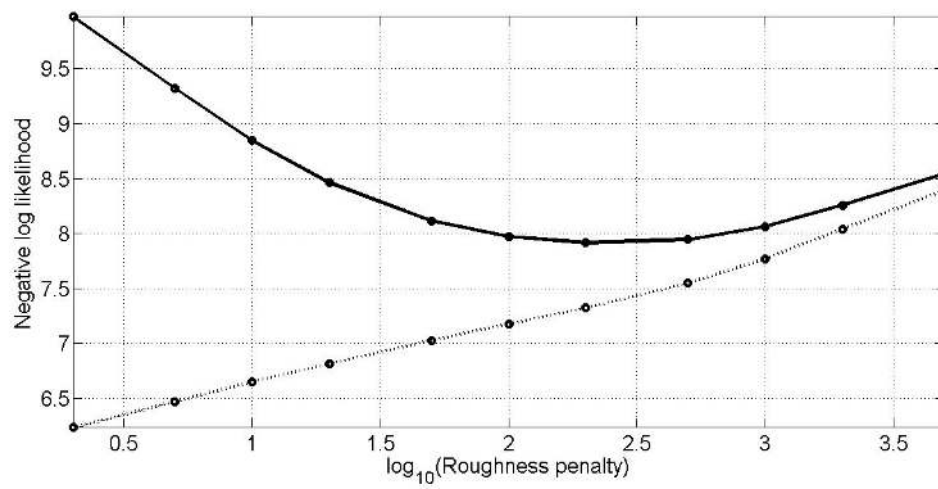


Figure 4. Simulation Case 1. Cross-validators selection of optimal roughness penalty λ for conditional extremes model. Negative log likelihood as a function of λ for cross-validation (black) and resubstitution (grey).

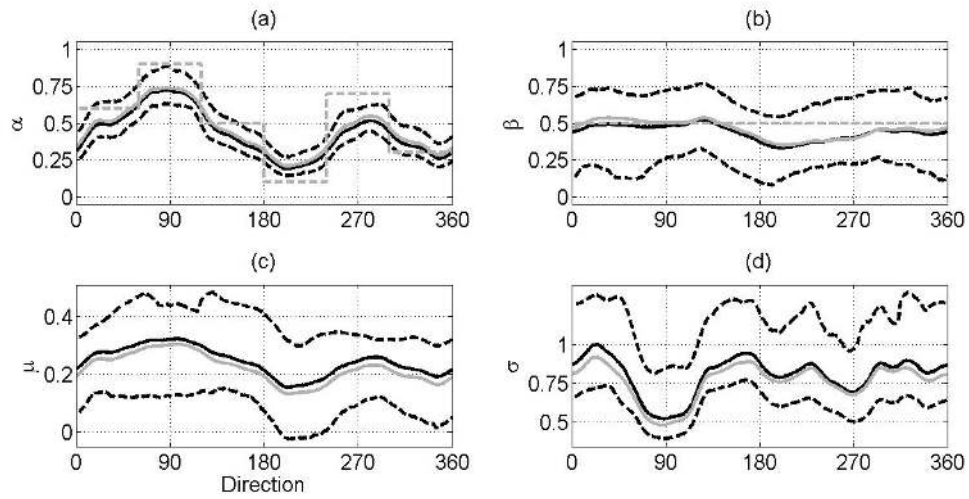


Figure 5. Simulation Case 1. Sample, bootstrap and true conditional extremes parameters with covariate. Sample estimates are given in solid grey. Median bootstrap estimates are given in solid black, with 95% bootstrap uncertainty bands in dashed black. True values of α and β in dashed grey.

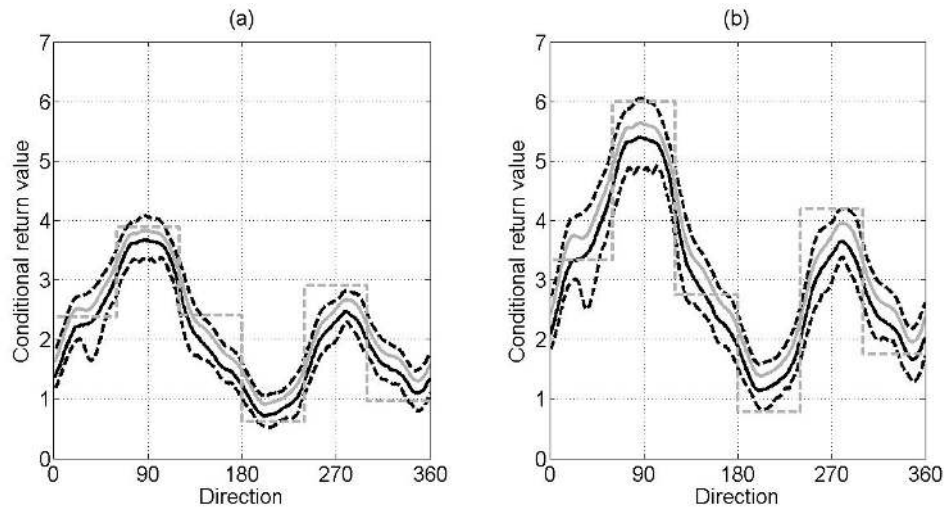


Figure 6. Simulation Case 1. Conditional return values of the associated variate \tilde{X}_2 , corresponding to values of the conditioning variate \tilde{X}_1 with non-exceedance probabilities (for period of sample) of (a) 0.99 and (b) 0.999. Bootstrap median (solid) and 95% uncertainty band (dashed) in black. Estimate using actual sample in solid grey. True values in dashed grey.

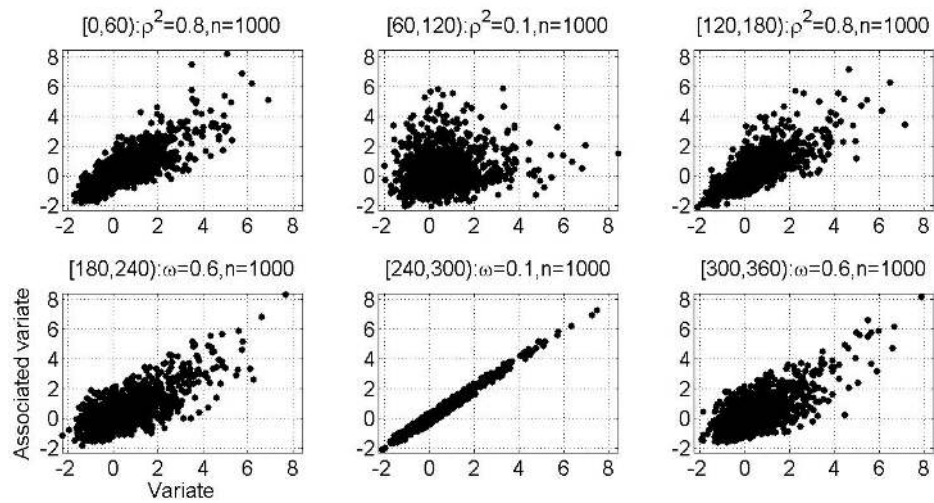


Figure 7. Simulation Case 2. Scatter plot per covariate interval. Values for intervals for covariate θ , parameters ρ^2, ω and sample size n are shown in each panel.

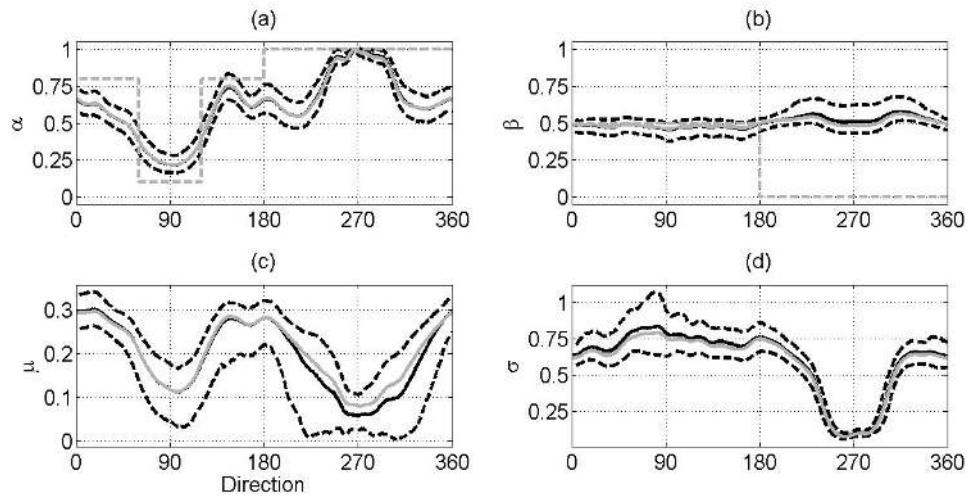


Figure 8. Simulation Case 2. Estimates for parameters α , β , μ and σ and their uncertainties as functions of covariate θ . Sample estimates are given in solid grey. Median bootstrap estimates are given in solid black, with 95% bootstrap uncertainty bands in dashed black. The true values of α and β are shown in dashed grey.




Figure 9. Simulation Case 2. Conditional return values of the associated variate \dot{X}_2 , corresponding to values of the conditioning variate \dot{X}_1 with non-exceedance probabilities (for period of sample) of (a) 0.99 and (b) 0.999. Bootstrap median (solid) and 95% uncertainty band (dashed) in black. Estimate using actual sample in solid grey. True values in dashed grey.




Figure 10. Northern North Sea. Non-crossing quantile estimates corresponding to non-exceedance probabilities of 0.1, 0.2, ..., 0.7, for (a) storm peak significant wave height and (b) associated storm peak period.

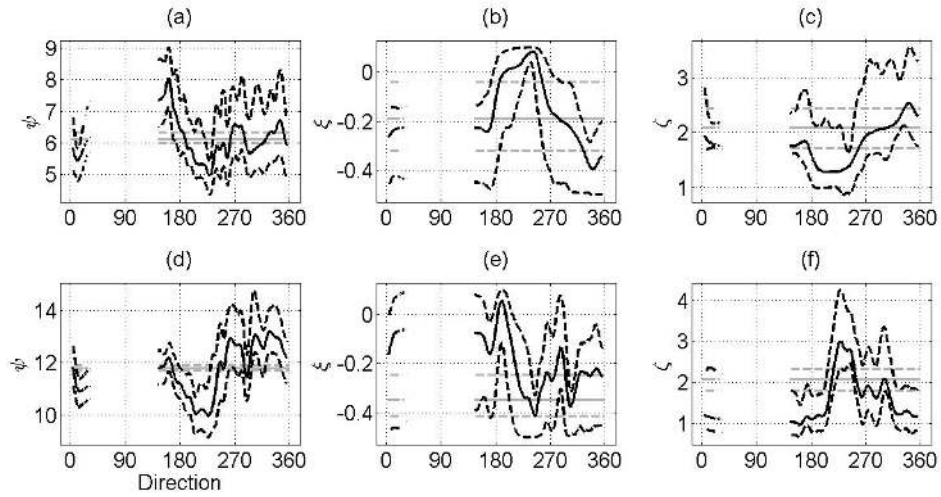


Figure 11. Northern North Sea. Sample (solid black) and 95% bootstrap uncertainty band (dashed black) estimates for parameters of marginal models for (a,b,c) extreme storm peak significant wave height and (d,e,f) associate storm peak period. (a,d) quantile regression threshold, (b,e) generalised Pareto shape, and (c,f) generalise Pareto scale. Corresponding estimates assuming no directional effect are given in grey.

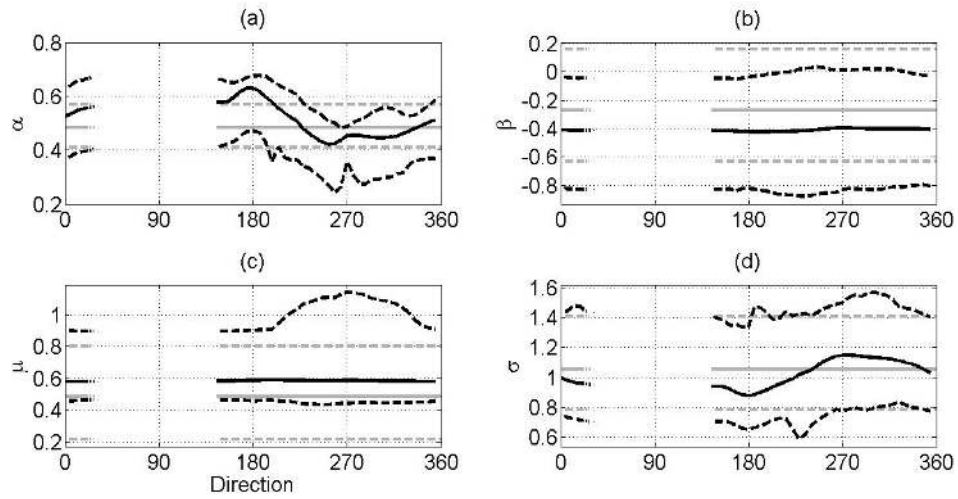


Figure 12. Northern North Sea. Non-stationary estimates for parameters α , β , μ and σ and their uncertainties (in black) as functions of covariate θ in terms of sample estimate (solid) and 95% bootstrap uncertainty bands (dashed). Corresponding stationary estimates in grey.

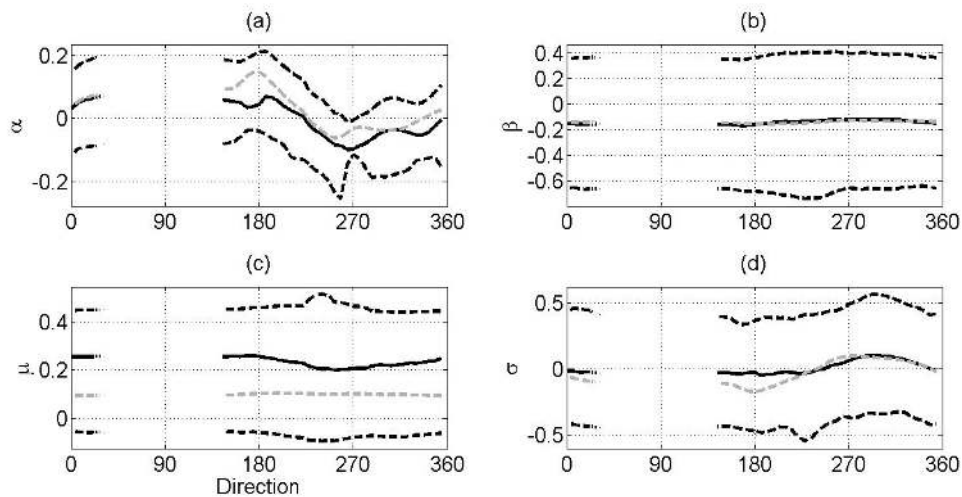


Figure 13. Northern North Sea. 95% bootstrap confidence intervals for the difference between non-stationary and stationary estimates for parameters α , β , μ and σ (in dashed black) as functions of covariate θ . Also shown are bootstrap median difference (solid black) and sample difference (grey)




Figure 14. Northern North Sea. Estimates for (a) marginal return values of associated storm peak period with non-exceedance probabilities (for period of sample) of 0.999, and (b) conditional return value of associated storm peak period given a value of storm peak significant wave height with non-exceedance probabilities (for period of sample) of 0.999. Estimates as functions of covariate θ (black) in terms of sample estimate (solid) and 95% bootstrap uncertainty band (dashed). Corresponding estimates assuming no directional dependence in grey.

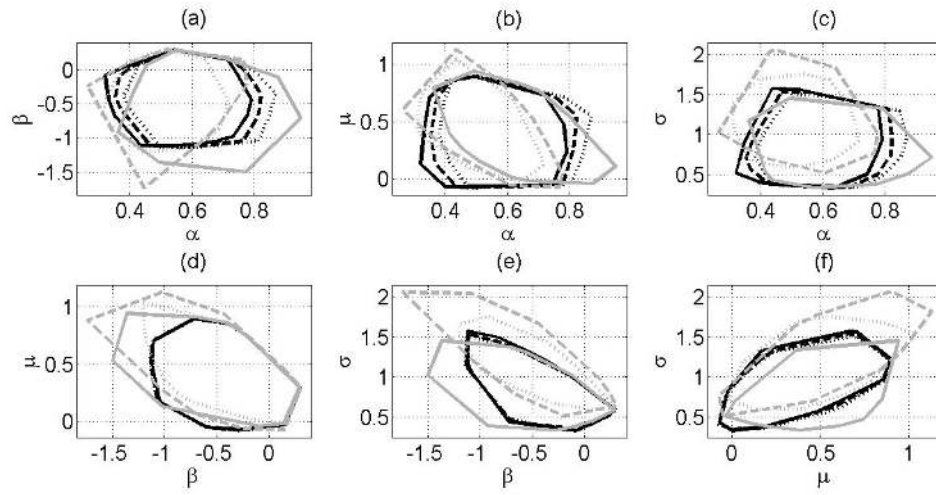


Figure 15. Northern North Sea. Convex hulls for pairs of parameter estimates corresponding to centre points of the directional sectors defined in Figure 2. Sectors 1 to 3 are in black, and 4 to 6 in grey. Sectors 1 and 4 are solid, 2 and 5 are dashed, 3 and 6 are dotted.