

# New estimation methods for extremal bivariate return curves

C. J. R. Murphy-Barltrop<sup>1\*</sup>, J. L. Wadsworth<sup>2</sup> and E. F. Eastoe<sup>2</sup>

<sup>1</sup>STOR-i Centre for Doctoral Training, Lancaster University LA1 4YR, United Kingdom

<sup>2</sup>Department of Mathematics and Statistics, Lancaster University LA1 4YF, United Kingdom

\*Correspondence to: c.barltrop@lancaster.ac.uk

February 16, 2023

## Abstract

In the multivariate setting, estimates of extremal risk measures are important in many contexts, such as environmental planning and structural engineering. In this paper, we propose new estimation methods for extremal bivariate return curves, a risk measure that is the natural bivariate extension to a return level. Unlike several existing techniques, our estimates are based on bivariate extreme value models that can capture both key forms of extremal dependence. We devise tools for validating return curve estimates, as well as representing their uncertainty, and compare a selection of curve estimation techniques through simulation studies. We apply the methodology to two metocean data sets, with diagnostics indicating generally good performance.

*Keywords:* Risk Measure, Extremes, Dependence Modelling

# 1 Introduction

## 1.1 Univariate extremal risk measures

Statistical analysis of extreme values is important in a wide range of environmental contexts, from the modelling of ocean waves (Shooter et al., 2021) and wind speeds (Baran et al., 2021) to climatological assessments (Yu et al., 2022; Worms and Naveau, 2022). The two most common approaches to modelling the extreme behaviour (or tail) of a single variable are block maxima and peaks over threshold (Coles, 2001). For the latter, which is more popular in practice, the generalised Pareto distribution (GPD) is used to model exceedances of some high threshold. This is justified through the Pickands-Balkema-de Haan theorem (Balkema and de Haan, 1974; Pickands, 1975); given a random variable  $X \sim F_X$  satisfying certain conditions, there exists a normalising function  $c(u)$  such that

$$\Pr\left(\frac{X-u}{c(u)} \leq x \mid X > u\right) \rightarrow G(x) := 1 - \left\{1 + \frac{\xi x}{\sigma}\right\}_+^{-1/\xi}, \quad x > 0, (\sigma, \xi) \in \mathbb{R}^+ \times \mathbb{R}, \quad (1.1)$$

as  $u \rightarrow x^F := \sup\{x : F_X(x) < 1\}$ . Here,  $G$  is the cumulative distribution function of a GPD, with scale and shape parameters,  $\sigma$  and  $\xi$ , respectively, and  $z_+ = \max(0, z)$ . The shape parameter  $\xi$  determines the behaviour of the tail: the cases  $\xi < 0$ ,  $\xi = 0$  and  $\xi > 0$  correspond to bounded, exponential and heavy tails, respectively. Given a sufficiently high threshold  $u$ , we assume the statistical model  $X - u \mid X > u \sim \text{GPD}(\sigma, \xi)$ . A detailed discussion of peaks over threshold modelling can be found in Coles (2001).

Often, univariate extreme value models are used to estimate risk measures for events associated with small probabilities; these summary statistics provide a quantification of extremal behaviour that can be used to help mitigate against rare environmental events, such as floods, storms, or wildfires. One such measure is known as a return level. Given a real variable  $X$ , representing a measurement taken at regular time intervals, and probability  $p$ , the  $p$ -probability return

level is the value  $x_p$  that satisfies the equation  $\Pr(X > x_p) = p$ . For small  $p$ ,  $x_p$  represents a high quantile which can be estimated using the GPD. We restrict attention to the case when  $X$  is stationary since the interpretation of return levels is more straightforward in this setting. We define the return period of  $x_p$  to be the value  $1/p$ ; one would expect the variable  $X$  to exceed  $x_p$  once, on average, during each return period. The relationship between return levels and periods can be illustrated using a return level plot; examples of three such plots with varying shape parameters are given in the left panel of Figure 1.




Figure 1: Left: Return level plots under varying shape parameters. Centre: Return curve  $RC(p)$  (red) of standard bivariate normal data with  $\rho = 0.5$  and  $n = 1000$  for return period of 100 ( $p = 1/100$ ). In a sample of size  $n$ , we expect to observe  $np$  points in the blue shaded region; this is true for all equivalent shaded regions at any point on the curve. Right: Return curves of the same bivariate normal data set for return periods in the set  $\{10, 100, 1000, 10000\}$ .

Return levels are widely used and provide a simple way to understand risk. However, many potentially impactful events arise due to the effect of more than one variable. For example, Mattei et al. (2001) describe how the combination of high sea levels and wind led to large-scale flooding at the Blayais nuclear power plant in 1999, causing significant damage. For this reason, it is desirable to have similar risk measures in the multivariate case, but thus far relatively little consideration has been given to this problem. This is in part due to the lack of natural ordering for multivariate vectors, which means there is no longer a single definition of an extreme event. A variety of measures have been proposed (Serinaldi, 2015), with each suited to a different analytical need. In this text, we focus on the bivariate case and one particular measure known as

a return curve, which directly extends the return level concept to the bivariate setting.

## 1.2 Return curves

Consider the joint survival function of the continuous random vector  $(X, Y)$  at a given probability  $p$ , i.e.  $\Pr(X > x, Y > y) = p$ . The combinations  $(x, y) \in \mathbb{R}^2$  satisfying this equation define a curve in the plane; we therefore define the  $p$ -probability return curve to be the set

$$\text{RC}(p) := \{(x, y) \in \mathbb{R}^2 : \Pr(X > x, Y > y) = p\}.$$

We consider values of  $p$  close to zero, corresponding to rare joint exceedance events. Within the literature, this set has a variety of labels, including isolines (Cooley et al., 2019), hazard curves (Simpson and Wadsworth, 2017) and joint probability curves (Gouldby et al., 2017). In an analogue to return levels, we define the return period to be  $1/p$ , since given any point  $(x, y) \in \text{RC}(p)$ , we would expect to observe the event  $\{X > x, Y > y\}$  once, on average, each return period. Equivalently, in a sample of size  $n$  from  $(X, Y)$ , we expect to observe  $np$  points in the region  $(x, \infty) \times (y, \infty)$ .

Since return curves define a line in  $\mathbb{R}^2$  rather than a single value, the two dimensional return level plot does not naturally extend to this setting. Instead, we can consider different return periods and plot the corresponding curves individually or simultaneously; examples of both are given in the centre and right panels of Figure 1 for a standard bivariate normal data set with correlation coefficient  $\rho = 0.5$ .

Return curves are arguably the most intuitive bivariate extension to return levels since they are also defined in terms of the survivor function. They have been used in practice to derive extremal environmental conditions for the design and analysis of ocean structures, such as oil rigs (Jonathan et al., 2014), freight ships (Vanem et al., 2020) and wind turbines (Manuel et al., 2018; Velarde et al., 2019), and coastal structures, such as railway lines (Environmental Agency,

2005; Gouldby et al., 2017) and wave energy converters (Eckert-Gallup and Martin, 2016).

As motivating examples, we consider two environmental data sets of practical importance, both of which are illustrated in Figure 2. Our objective is to use return curve estimates to derive joint extremal conditions for each data set. Both data sets are comprised of metocean variables, which have previously been used in a comparison exercise for a risk measure known as an environmental contour (Haselsteiner et al., 2021); such measures also aim to summarise joint extremal behaviour. However, unlike return curves, they do not offer an intuitive interpretation in terms of return periods.

The first data set contains measured significant wave height (m) and zero up crossing period (s) between 1996-2005 obtained from a buoy on the east coast of Florida, USA. The second data set contains 25 years of wind speed (m/s) and significant wave height (m) observations obtained from the hindcast model coastDat -2 (Groll and Weisse, 2017) for a location in the North Sea near the east coast of the UK. These combinations of variables are of particular relevance for the structural reliability of offshore and coastal structures, and bivariate risk measures are commonly used to inform the design basis for such structures (Jonathan et al., 2014; Haselsteiner et al., 2019; Mackay and Jonathan, 2020). They therefore provide realistic examples with which to illustrate the utility of return curve estimates.

Both of the original data sets are comprised of hourly observations, resulting in strong temporal dependence between observations. To reduce the temporal dependence while retaining a relevant data set, we restrict attention to daily maxima values. Furthermore, we account for seasonal non-stationarity by restricting attention to the months September-March; our analysis suggests the largest wave heights are observed in this interval. This results in  $n = 2048$  and  $n = 5306$  observations, respectively, for the measured and hindcast data sets.




Figure 2: Left: Measured metocean data for a buoy located on the east coast of Florida, USA. Right: Hindcast metocean data for a location near the east coast of the UK.

### 1.3 Practical considerations for return curve estimation

To estimate return curves in practice, we must first estimate the joint survival function of an observed bivariate process at a fixed probability  $p$ . Since we restrict attention to probabilities close to zero, we require an estimation method that can accurately capture the joint tail behaviour of the process and provide realistic extrapolation to estimate  $RC(p)$  for  $p$  outside of the observation period. For this reason, we restrict attention to models from multivariate extreme value theory. We must also consider methods for evaluating the uncertainty associated with return curve estimates. This is a more complex problem than assessing the uncertainty of return level estimates, since there is more than one dimension in which the risk measure can vary. Furthermore, given a data set for which a curve has been estimated at probability  $p$ , it is also essential to evaluate goodness of fit via a diagnostic tool. Little attention has been given to either of these problems within the extremes literature, motivating the development of novel approaches.

This article is structured as follows. Section 2 provides a brief background on bivariate extreme value theory. In Section 3, we introduce various properties of return curves, review existing

approaches for curve estimation within the extremes literature and present two novel approaches. Section 4 introduces new tools for analysing uncertainty and goodness of fit in return curve estimates. Section 5 presents a simulation study to evaluate the performance of these tools and compare curve estimates from several models. In Section 6, we apply our methodology to the two aforementioned environmental data sets. We conclude in Section 7 with a discussion and outlook on future work.

## 2 Bivariate extreme value theory

When assessing the extremal behaviour of a continuous bivariate random vector  $(X, Y)$ , one must consider tail behaviour within both marginal processes, along with the dependence between the largest observations of either variable, which we refer to as the extremal dependence. A fundamental classification of extremal dependence is whether the most extreme events can occur jointly or not. This is quantified by the coefficient  $\chi := \lim_{u \rightarrow 1} \chi(u) \in [0, 1]$ , where

$$\chi(u) = \Pr(F_Y(Y) > u \mid F_X(X) > u),$$

for  $X \sim F_X$ ,  $Y \sim F_Y$ . The cases  $\chi = 0$  and  $\chi > 0$  correspond to asymptotic independence (AI) and asymptotic dependence (AD), respectively. A larger  $\chi$  corresponds to stronger positive dependence in the joint tail, with  $\chi = 1$  occurring for perfectly dependent variables. In practice, we cannot estimate  $\chi$  in the limit as  $u \rightarrow 1$ , meaning we must use  $\chi(u)$ ,  $u < 1$ , to determine AD vs AI.

Early modelling techniques proposed for multivariate extreme values focused only on the AD case (de Haan and Resnick, 1977; Coles and Tawn, 1991; de Haan and de Ronde, 1998). Such methods are based on the framework of multivariate regular variation: given a random vector  $(X, Y)$  with standard Fréchet margins, we define radial and angular components  $R := X + Y$

and  $W := X/R$ , respectively. We say  $(X, Y)$  is multivariate regularly varying if, for measurable  $B \subset [0, 1]$ , we have

$$\lim_{r \rightarrow \infty} \Pr(W \in B, R > sr \mid R > r) = H(B)s^{-1}, \quad s \geq 1, \quad (2.1)$$

with  $H(\partial B) = 0$ , where  $\partial B$  is the boundary of  $B$  (Resnick, 1987). Assumption (2.1) implies that for large radial values,  $R$  and  $W$  are independent. The quantity  $H$ , which is known as the spectral measure, captures the extremal dependence structure of  $(X, Y)$  and must satisfy the moment constraint  $\int_0^1 w dH(w) = 1/2$ . For all AI vectors, the spectral measure places mass on the atoms  $\{0\}$  and  $\{1\}$ ; as such, this modelling framework cannot capture tail properties under this extremal dependence scheme (Coles et al., 1999).

In recent years, it has been shown that the AI case is at least as important as the AD case, and that assuming the incorrect form of extremal dependence leads to unsuitable extrapolation in the joint upper tail (Ledford and Tawn, 1996, 1997; Heffernan and Tawn, 2004). Therefore, unless there is strong prior knowledge in favour of either AD or AI, it is desirable to use models which have sufficient flexibility to allow the data to directly inform the class of extremal dependence structure.

The first such approach was proposed in Ledford and Tawn (1996). Given  $(X, Y)$  with standard exponential margins, this model assumes the joint tail representation

$$\Pr(X > t, Y > t) = \Pr(\min(X, Y) > t) = L(e^t)e^{-t/\eta} \quad \text{as } t \rightarrow \infty, \quad (2.2)$$

where  $L$  is a slowly varying function at infinity, i.e.,  $\lim_{t \rightarrow \infty} L(ct)/L(t) = 1$  for  $c > 0$ , and  $\eta \in (0, 1]$ . The parameter  $\eta$  is termed the coefficient of tail dependence, with  $\eta = 1$  and  $\lim_{t \rightarrow \infty} L(t) > 0$  corresponding to AD and  $\eta < 1$ , or  $\eta = 1$  and  $\lim_{t \rightarrow \infty} L(t) = 0$ , corresponding to AI. Estimation of  $\eta$  can be performed in practice using the Hill estimator (Hill, 1975). Several extensions to this



approach exist (Ledford and Tawn, 1997; Ramos and Ledford, 2009): however, all techniques derived under this framework are applicable only within regions where both variables are large. Consequently, these methods are not appropriate for the estimation of return curves, since this measure is defined also in regions where only one variable is large; see Figure 1.

Wadsworth and Tawn (2013) provide an alternative representation for bivariate tail probabilities using a more general extension of the model described in equation (2.2) that allows for joint tail estimation in regions where only one variable is large. Given  $(X, Y)$  with standard exponential margins, they assume for each  $w \in [0, 1]$

$$\Pr(\min\{X/w, Y/(1-w)\} > t) = L(e^t | w)e^{-\lambda(w)t}, \quad \lambda(w) \geq \max(w, 1-w), \quad (2.3)$$

as  $t \rightarrow \infty$ , where  $L(\cdot | w)$  is slowly varying for each ray  $w \in [0, 1]$ . The function  $\lambda$ , which is termed the angular dependence function, is the key quantity in determining joint tail behaviour, and both AD and AI can be captured under this assumption, with AD implying the lower bound  $\lambda(w) = \max(w, 1-w)$ . This quantity generalises the coefficient  $\eta$ , with  $\eta = 1/(2\lambda(0.5))$ , and can be estimated pointwise for any ray  $w \in [0, 1]$  using the Hill estimator. This approach can be used to estimate joint survivor probabilities where only one variable is large by taking values of  $w$  close to 0 or 1.

Heffernan and Tawn (2004) proposed a very general modelling tool for conditional probabilities. We consider the extension given in Keef et al. (2013) since the formulation given in the original approach cannot easily accommodate structures exhibiting negative dependence. Given a random vector  $(X, Y)$  with standard Laplace margins, it is assumed that there exist normalising functions  $a : \mathbb{R} \rightarrow \mathbb{R}$  and  $b : \mathbb{R} \rightarrow \mathbb{R}_+$  such that

$$\lim_{t \rightarrow \infty} \Pr[(Y - a(X))/b(X) \leq z, X - t > x | X > t] = D(z)e^{-x}, \quad (2.4)$$

for a non-degenerate distribution function  $D$ . Similarly to Wadsworth and Tawn (2013), this framework is able to capture both AD and AI, with AD arising when  $a(x) = x$  and  $b(x) = 1$ . This method is a flexible approach for modelling multivariate extremes and is also not restricted only to regions where both variables are large. Note that one could also condition on the event  $Y > t$  and assume the existence of normalising functions for the variable  $X$ : in combination, these assumptions allow consideration of the region where either variable is large. The functions  $a$  and  $b$  are typically estimated parametrically under a misspecified model for  $D$ , while the distribution function  $D$  is subsequently estimated non-parametrically.

Alongside these approaches, we note that there exist a range of copula-based models that can capture both dependence regimes (Coles and Pauli, 2002; Wadsworth et al., 2017; Huser and Wadsworth, 2019). Such techniques aim to create a unified modelling framework for AD and AI. Moreover, the case of AD does not represent a boundary case for the latter two approaches, which could be practically advantageous. However, they all require stronger assumptions about the form of parametric family for the bivariate distribution, reducing their flexibility and limiting their use in practice. As a result, we prefer instead to consider the more flexible modelling assumptions described in equations (2.3) and (2.4).

### **3 Bivariate return curve estimation**

We now consider techniques for practical estimation of return curves. We begin by introducing theoretical results for return curves in Section 3.1. Naive implementation of statistical estimation methods would generally produce curves that fail to respect these results, but imposing them will typically improve estimation. Section 3.2 details how to perform marginal transformations so that the models in equations (2.3) and (2.4) can be applied in practice. In Section 3.3, we present existing approaches for return curve estimation before introducing novel estimation techniques in Section 3.4.

### 3.1 Return curve properties

A careful consideration of the theory surrounding the joint survival function allows us to deduce several properties about the shape and magnitude of  $\text{RC}(p)$  for a given  $p \in (0, 1)$ . We begin by noting that the joint distribution function,  $F_{X,Y}$ , can be expressed in terms of the marginal distribution functions of  $X$  and  $Y$ ,  $F_X, F_Y$ , respectively, and a copula  $C$  via  $F_{X,Y}(x, y) = C(F_X(x), F_Y(y))$ . The return curve is linked to the joint distribution function  $F_{X,Y}$  by the equation  $\Pr(X > x, Y > y) = 1 - F_X(x) - F_Y(y) + F_{X,Y}(x, y)$ . Throughout this section, we make the assumption that the random vector  $(X, Y)$  has strictly continuous marginal distribution functions.

**Property 3.1.** *Let  $x_p := F_X^{-1}(1 - p)$  and  $y_p := F_Y^{-1}(1 - p)$  be the  $(1 - p)$ -th quantiles of  $X$  and  $Y$ , respectively. Then for  $(x, y) \in \text{RC}(p)$ ,  $x \leq x_p$  and  $y \leq y_p$ .*

*Proof.* We have  $\Pr(X > x) \geq \Pr(X > x, Y > y) = p = \Pr(X > x_p)$  and hence  $x \leq x_p$ .

□

This result bounds the coordinate values that can be observed on the return curve. Next, by considering the limit of the joint survivor function as one variable converges to the lower limit of the marginal support, we obtain the following result.

**Property 3.2.** *Let  $\text{supp}(F)$  denote the support of  $F$  and  $x_{inf} := \inf\{\text{supp}(F_X)\}$ ,  $y_{inf} := \inf\{\text{supp}(F_Y)\}$ .*

*We have that*

$$\Pr(X > x, Y > y) = \begin{cases} \Pr(Y > y) & \text{if } x \leq x_{inf}, \\ \Pr(X > x) & \text{if } y \leq y_{inf}. \end{cases}$$

Combining this statement with Property 3.1, Property 3.3 follows.

**Property 3.3.** *Let  $(x, y) \in \text{RC}(p)$ . If  $x \leq x_{inf}$  ( $y \leq y_{inf}$ ), then  $y = y_p$  ( $x = x_p$ ).*

These results allow us to easily compute the curve coordinates on the regions  $(-\infty, x_{inf}) \times (y_p, \infty)$  and  $(x_p, \infty) \times (-\infty, y_{inf})$ , assuming we can accurately estimate the marginal quantiles

$(x_p, y_p)$  and the infima of marginal supports  $(x_{inf}, y_{inf})$ . Finally, by considering coordinates at different points on a return curve, we obtain the following result.

**Property 3.4.** *Suppose the copula,  $C$ , of  $(X, Y)$  on uniform margins has joint support on the whole of  $[0, 1]^2$  and joint density function, denoted  $c$ . Given  $(x_1, y_1), (x_2, y_2) \in \text{RC}(p)$  with  $0 < F_X(x_1), F_X(x_2) < 1$  and  $0 < F_Y(y_1), F_Y(y_2) < 1$ , we have that  $x_1 < x_2 \Leftrightarrow y_1 > y_2$ .*

*Proof.* Suppose  $x_1 < x_2$  and  $y_1 \leq y_2$ . This implies that

$$\begin{aligned}
p &= \Pr(X > x_1, Y > y_1) = \Pr(F_X(X) > F_X(x_1), F_Y(Y) > F_Y(y_1)) \\
&= \int_{F_X(x_1)}^1 \int_{F_Y(y_1)}^1 c(u, v) \, dv \, du \\
&= \int_{F_X(x_1)}^{F_X(x_2)} \int_{F_Y(y_1)}^1 c(u, v) \, dv \, du + \int_{F_X(x_2)}^1 \int_{F_Y(y_1)}^1 c(u, v) \, dv \, du \\
&> \int_{F_X(x_2)}^1 \int_{F_Y(y_1)}^1 c(u, v) \, dv \, du \quad (\text{since we have support on the whole of } [0, 1]^2) \\
&\geq \int_{F_X(x_2)}^1 \int_{F_Y(y_2)}^1 c(u, v) \, dv \, du = \Pr(X > x_2, Y > y_2) = p
\end{aligned}$$

implying  $p > p$ , a contradiction. Hence,  $y_1 > y_2$ . □

This result governs the shape of the contour defined by the return curve set. We note there is an alternative proof given in Cooley et al. (2019) under the assumption of monotonicity of the joint survivor function.

## 3.2 Marginal transformations

From Section 2, it is clear that in order to apply multivariate extreme value models, we need to standardise the marginal distributions of a random vector to achieve the form assumed by the model. Typically, inference involves two steps: forward transformation to get the data onto desired margins and back transformation to move any computed statistics, such as a return curve, back onto the original margins. For both steps, we use the semi-parametric approach given in

Coles and Tawn (1991). Given an identically distributed sample  $\{(x_i, y_i) : i = 1, \dots, n\}$  from a random vector  $(X, Y)$  with unknown margins, we estimate the marginal distribution  $\hat{F}_X$  (similarly  $\hat{F}_Y$ ) by

$$\hat{F}_X(x) = \begin{cases} 1 - \{1 - \tilde{F}_X(u_X)\} \{1 + \xi_X(x - u_X) / \sigma_X\}_+^{-1/\xi_X}, & \text{for } x > u_X, \\ \tilde{F}_X(x), & \text{for } x \leq u_X, \end{cases} \quad (3.1)$$

where the first line represents the GPD above a high threshold  $u_X$  and  $\tilde{F}_X$  is an empirical rank transform given by  $\tilde{F}_X(x) = \sum_{i=1}^n \mathbb{1}(x_i \leq x) / (n + 1)$ . This approach ensures the marginal tail behaviour is captured within the transformation. Moreover, equation (3.1) can be easily inverted to perform the back transformation step.

### 3.3 Existing methodology

The literature on extremal return curve estimation is sparse, owing to the fact that little consideration has been given to this problem in practice. Of the existing approaches, each can be designated into one of three categories: approaches for AD data only, approaches for AI data only, and approaches applicable to data exhibiting either regime. The majority of the available literature falls within the first of these categories (Stephenson, 2002; Salvadori and De Michele, 2004; Marcon et al., 2017). In all cases, the authors assume the bivariate copula is in a family of distributions termed bivariate extreme value copulas. These copulas, which are directly related to the spectral measure in equation (2.1), imply AD and provide the basis for the majority of multivariate extreme value techniques. The necessity for making such a strict assumption is a well known drawback of this kind of model, since the form of extremal dependence is seldom known prior to analysis and AI is frequently observed in practice (Heffernan and Tawn, 2004; Huser and Wadsworth, 2019). We therefore choose not to consider such approaches further.

In Cooley et al. (2019), the authors propose separate techniques for the first and second categories. In both cases, extremal return curve estimates are obtained by ‘projecting’ empirical

curves estimated for less extreme probabilities. For the case of AD, given a random vector  $(X, Y)$  with standard Fréchet margins, multivariate regular variation is exploited to obtain curve estimates. Consider two small probabilities  $p$  and  $p^*$  with  $p^* > p$ ; multivariate regular variation implies that  $\text{RC}(p) \approx s^{-1}\text{RC}(p^*)$ , where  $s := p^*/p > 1$ . In practice,  $\text{RC}(p^*)$  is estimated empirically via a smooth Gaussian-kernel estimate of the joint survivor function, and scaled by the coefficient  $s^{-1}$  to produce an estimate for  $\text{RC}(p)$ .

In the case of AI, a similar estimation procedure is proposed based on the framework of hidden regular variation (Resnick, 2002), an elaboration of the assumption outlined in equation (2.2). However, as mentioned in Section 2, this approach only works in regions where both variables are large. To account for this, Cooley et al. (2019) proposed an ad-hoc procedure to link this region to the marginal axes, requiring additional steps and parameter estimation. In contrast, the assumption in equation (2.3) provides a theoretically sound link between regions where variables are of different magnitudes. Estimates of return curves from the Cooley et al. (2019) approach are illustrated in Figure 8. In Section 5, we compare the resulting curve estimates from this approach to the methods introduced in this paper and show that the techniques we present outperform this method in a wide range of scenarios.

For the third category, few approaches exist within the literature; this is in part because the bivariate extreme value methodologies that allow for dual estimation are relatively modern. All proposed techniques use a semi-parametric implementation of the conditional extremes model described in equation (2.4) (Jonathan et al., 2014; Gouldby et al., 2017; Simpson and Wadsworth, 2017; Tilloy et al., 2020). However, like the other techniques introduced in this section, little to no consideration is given to the theory behind return curves, leading to curve estimates with undesirable properties. Moreover, as will be discussed in Section 3.4.1, utilising the Heffernan and Tawn (2004) modelling framework for return curve estimation is not straightforward, requiring delicate treatment and several steps; this has not been fully acknowledged in these existing approaches.

To the best of our knowledge, the modelling techniques discussed here cover all of the proposed methods for estimating return curves at extremal probabilities. Furthermore, we know of no attempt to compare curve estimates from these different methods. No formal quantifications of return curve uncertainty or bias have been proposed previously, making it difficult to evaluate performance over different dependence structures. Some approaches (Simpson and Wadsworth, 2017; Cooley et al., 2019) instead provide bootstrap curve estimates which, while representing the uncertainty in curve estimates, do not provide interpretable confidence regions.

Alongside this issue, there is only one diagnostic tool in the literature for evaluating the accuracy of return curve estimates (Cooley et al., 2019). This tool utilised the result that, if data are independent and identically distributed, the number of points in each survival region on the return curve should theoretically be  $\text{Binomial}(n, p)$  distributed, where  $n$  denotes the size of the data set. This property can be used to construct a theoretical confidence region for the number of points in the joint survival set  $(x, \infty) \times (y, \infty)$  for any point  $(x, y) \in \text{RC}(p)$ . While the authors show that these confidence regions capture the number of observations within estimated joint survival sets for the majority of considered examples, we argue that the resulting diagnostic is relatively uninformative since it lacks an intuitive interpretation in terms of the survival probability  $p$ . Moreover, this diagnostic strongly relies on the assumption of independent observations, which is seldom the case in practice. We present an alternative diagnostic tool in Section 4.2, where confidence intervals are instead obtained using sets of empirical probability estimates obtained through bootstrapping and compared to the true probability  $p$ . Temporal dependence can be incorporated through block bootstrapping, meaning this tool can be applied to a wider range of data sets; see Sections 4.2 and 6 for further details.

### **3.4 Novel methods for return curve estimation**

We outline two methods for estimation of  $\text{RC}(p)$  based on the modelling assumptions given in equations (2.3) and (2.4). Consider a random vector  $(X, Y)$  with standard exponential margins,

for which the marginal support is given by the set  $\mathbb{R}_+$ ; this implies  $x_{inf} = y_{inf} = 0$ . We can immediately deduce from Property 3.3 that the coordinates of the return curve intersecting the margins are given by  $(0, y_p)$  and  $(x_p, 0)$ , with  $y_p = x_p = -\log(p)$ , the  $(1 - p)$ -th quantile. These coordinates give us ‘start’ and ‘end’ points for curve construction. Moreover, given a curve estimate  $\widehat{RC}(p)$  for this vector constructed with these boundary points, Properties 3.1 and 3.4 can be imposed to ensure the resulting curve is theoretically possible. For the former, if any  $(x, y) \in \widehat{RC}(p)$  satisfy  $x > x_p$  (similarly,  $y > y_p$ ), we set  $x = x_p$  ( $y = y_p$ ), thereby bounding values on the curve. For the latter, we treat the bounded curve estimate as a function of  $x$  and apply an iterative algorithm starting at the point  $(0, y_p)$  to obtain a monotonic function. Incorporating additional theoretical knowledge into return curve estimation should lead to more accurate and robust estimates. For both of the methods introduced in this section, we impose the properties above retrospectively once curve estimates have been obtained.

### 3.4.1 Method based on the approach given in Heffernan and Tawn (2004)

In this section, we propose an implementation of the Heffernan and Tawn (2004) model, which builds on the existing methods that have applied this framework for return curve construction. Unlike these techniques, we incorporate the properties introduced in Section 3.1 into return curve estimates and provide an intuitive algorithm for combining the point estimates obtained from conditioning on both variables. Let  $(X_L, Y_L)$  denote the vector  $(X, Y)$  on standard Laplace margins and consider a small probability  $p$  for which we wish to obtain a return curve estimate. To achieve this, we fit the Heffernan and Tawn (2004) model twice, conditioning on both  $X_L$  and  $Y_L$  separately, thus allowing us to estimate the curve in different regions. In particular, we consider the regions defined by  $R_{Y_L > X_L} := \{(x_L, y_L) \in \mathbb{R}^2 \mid y_L > x_L\}$  and  $R_{Y_L \leq X_L} := \{(x_L, y_L) \in \mathbb{R}^2 \mid y_L \leq x_L\}$ .

For  $R_{Y_L > X_L}$ , we first select a high quantile  $u_{Y_L}$  from the marginal distribution of  $Y_L$  such that  $\Pr(Y_L > u_{Y_L}) > p$ . In particular, we select the 0.95 empirical quantile of this distribution, implying the return curve probability,  $p$ , must be smaller than 0.05. We assume the normalising functions



are given by  $a(y) = \alpha y$  and  $b(y) = y^\beta$  for constants  $\alpha \in [-1, 1]$  and  $\beta \in (-\infty, 1)$ : as noted in Keef et al. (2013), these functions capture the limiting dependence structures for a wide range of distributions. The parameters  $\alpha$  and  $\beta$  can be estimated under the working assumption that the distribution function  $D$ , which captures the stochastic behaviour of the variable  $(X_L - \alpha Y_L)/Y_L^\beta \mid Y_L > u_{Y_L}$ , is that of a Gaussian distribution. We denote the fitted values by  $\hat{\alpha}$  and  $\hat{\beta}$ . These values can be used to simulate from the variable  $X_L \mid Y_L > u_{Y_L}$ ; for example, see Jonathan et al. (2014).

We then consider a decreasing sequence of high quantiles from  $Y_L$  that exist in the interval  $(u_{Y_L}, F_{Y_L}^{-1}(1 - p))$ . The upper end point of this interval is the limit that  $Y_L$  values can attain on this curve and the lower end point represents the minimal quantile for which the fitted model is valid. We denote this set by  $\mathcal{Y}$  and iteratively consider each  $y_* \in \mathcal{Y}$  in turn, with  $q := \Pr(Y_L > y_*)$ . Using the fitted parameter values, we use the model to simulate from the conditional distribution  $X_L \mid Y_L > y_*$ . Letting  $x_*$  denote the (estimated)  $(1 - p/q)$ -th quantile from this distribution, we have that the resulting coordinate  $(x_*, y_*)$  is a member of the set  $\widehat{RC}(p)$  (defined for  $(X_L, Y_L)$ ) since  $\Pr(X_L > x_*, Y_L > y_*) = \Pr(X_L > x_* \mid Y_L > y_*)\Pr(Y_L > y_*) = \frac{p}{q} \times q = p$ . We continue in this manner until we obtain a value  $x_{**}$  with  $y_* \leq x_{**}$ , or we have exhausted all values in the set  $\mathcal{Y}$ . The resulting coordinate set then gives an estimate of the curve in  $R_{Y_L > X_L}$ .

A near identical procedure is used to obtain the curve estimate in  $R_{Y_L \leq X_L}$ , this time selecting a high quantile  $u_{X_L}$  from the distribution of  $X_L$  and fitting the conditional model above this quantile. We then consider a set of quantiles in the interval  $(x'_{**}, F_{X_L}^{-1}(1 - p))$ , where  $x'_{**} = x_{**}$  if  $x_{**}$  exists and  $u_{X_L}$  otherwise, where  $u_{X_L}$  denotes the empirical 0.95 quantile from  $X_L$ . We label this set  $\mathcal{X}$ , ordered such that the quantiles are increasing, and use the fitted model to obtain quantiles from the conditional distribution  $Y_L \mid X_L > x_*$  for each  $x_* \in \mathcal{X}$ . The resulting coordinate sets from both regions are combined to give an estimate of the return curve over the entire joint support of  $(X_L, Y_L)$ . An illustration of this procedure is given in Figure 3. As can be observed, two sets of points estimates are obtained by conditioning on either variable; these sets are then combined to give an estimate of the entire return curve. As a final step, we apply the probability integral

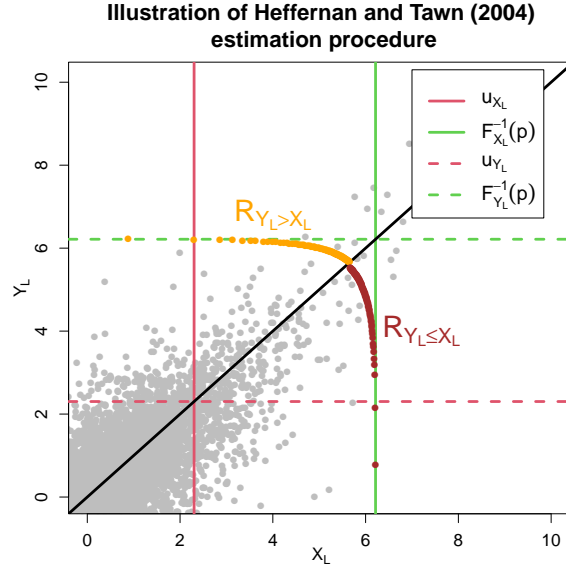


Figure 3: Illustration of return curve estimation procedure using Heffernan and Tawn (2004) modelling framework. Red and green lines give 0.95 and  $(1 - p)$ -th quantiles for both margins, respectively. The brown and orange points denote the point estimates obtained for the regions  $R_{Y_L > X_L}$  and  $R_{Y_L < X_L}$ , respectively.

transform to transform the curve estimate to standard exponential margins.

We note that the implementation of this model to estimate return curves is more complex than the methods proposed in Section 3.4.2 and Cooley et al. (2019). This is due to the fact the model requires a variable to condition on, meaning we have to fit the model twice and provide a technique for joining point estimates from regions  $R_{Y_L > X_L}$  and  $R_{Y_L < X_L}$ .

### 3.4.2 Method based on the approach given in Wadsworth and Tawn (2013)

In this section, we propose a novel implementation of the model described in equation (2.3) to generate non-parametric return curve estimates. Consider a random vector  $(X, Y)$  with standard exponential margins and define a set  $\mathscr{W}$  containing equally spaced rays in the interval  $[0, 1]$ , ordered from lowest to highest. Assuming  $|\mathscr{W}|$  is sufficiently large, we are able to evaluate the joint extremal behaviour across the entire region for which at least one variable is extreme. For each  $w \in \mathscr{W}$ , we use the 95% empirical threshold of the variable  $T_w := \min\{X/w, Y/(1 - w)\}$  to

obtain an estimate of the angular dependence function via the Hill estimator, which we denote  $\hat{\lambda}(w)$ . For large  $u$ , equation (2.3) implies that for any  $w \in (0, 1)$  and  $t > 0$ ,

$$\Pr\left(T_w > t + u \mid T_w > u\right) \approx \exp\{-t\hat{\lambda}(w)\}.$$

Estimates of  $t$  and  $u$ , combined with the rays  $w$ , provide estimates of points in  $\text{RC}(p)$ . We firstly select a small probability  $p^* > p$  and estimate  $u$  as the  $(1 - p^*)$ -th quantile of  $T_w$ , implying  $\Pr(T_w > u) = p^*$ . One can then estimate the value of  $t > 0$  such that  $\Pr(T_w > t + u) = p$  since

$$p = \Pr(T_w > t + u) = \Pr(T_w > u) \times \Pr(T_w > t + u \mid T_w > u) = p^* \exp\{-t\hat{\lambda}(w)\},$$

giving the estimate  $t = -\frac{1}{\hat{\lambda}(w)} \log(p/p^*)$ . Setting  $(x, y) := (w(t + u), (1 - w)(t + u))$ , we have  $(x, y) \in \widehat{\text{RC}}(p)$ , resulting in a return curve point estimate for each ray  $w \in \mathcal{W}$ .

## 4 Uncertainty estimation and diagnostic tool

### 4.1 Capturing uncertainty in return curve estimates

While previous methods for return curve estimation have considered sampling uncertainty, none provide a means to construct interpretable confidence regions and/or ‘average’ estimates for return curves (Simpson and Wadsworth, 2017; Cooley et al., 2019). Here, we propose a new method for representing uncertainty in return curve estimates that addresses limitations in the existing methods and provides a formal framework for comparing curve estimates from different models where the truth is known.

Our goal is to represent sampling uncertainty in return curve estimates via some type of confidence region at a given significance level  $\alpha \in (0, 1)$ . Since these curves vary in two dimensions, careful consideration is needed to ensure the resulting region represents  $\alpha$  in a straightforward

and interpretable manner.

Figure 4 displays  $n = 10000$  datapoints from inverted logistic (Ledford and Tawn, 1997) and asymmetric logistic (Tawn, 1988) copulas on standard exponential margins. The true return curves for  $p = 1/10000$  are given in red while the curves estimated using the Wadsworth and Tawn (2013) model are given in green. A representation of sampling uncertainty will help to determine the quality of these estimates. To achieve this, we propose an adaptation of a tool given in Haselsteiner et al. (2019) for representing uncertainty in environmental contour estimates. The novelty in our approach comes from the fact the original tool has not been directly applied for return curve estimation, even though return curves and environmental contours bear many similarities (Haselsteiner et al., 2021). Moreover, no consideration is given to the theoretical justification of the resulting uncertainty representation in the original approach in terms of coverage properties.




Figure 4: Estimated (green) vs true (red) curves for data sets simulated from inverted logistic (left) and asymmetric logistic (right) copulas.

On standard exponential margins, the joint support is given by the set  $\mathbb{R}_+^2$ ; we consider a set of angles in the interval  $(0, \pi/2)$  and use these angles to divide the  $\mathbb{R}_+^2$  plane. Specifically, given some large positive integer  $m$ , we define  $\Theta := \{\pi(m+1-j)/2(m+1) \mid 1 \leq j \leq m\}$ , i.e., a sequence of decreasing angles starting near  $\pi/2$  and approaching 0. For each  $\theta \in \Theta$ , let

$L_\theta := \{(x, y) \in \mathbb{R}_+^2 \mid \tan(\theta) = y/x\}$  denote the line segment intersecting the origin with gradient  $\tan(\theta) > 0$ . For any return curve estimate  $\widehat{\text{RC}}(p)$  satisfying properties 3.1 - 3.4, we have that  $L_\theta$  intersects  $\widehat{\text{RC}}(p)$  exactly once for every  $\theta \in \Theta$ , implying there is a one-to-one correspondence between angles and points on the estimated curve. An illustrative figure of this correspondence can be found in the supplementary material.

Letting  $\{(\hat{x}_\theta, \hat{y}_\theta)\} := \widehat{\text{RC}}(p) \cap L_\theta$ , we let  $\hat{d}_\theta$  denote the  $l_2$ -norm of this point estimate, i.e.,  $\hat{d}_\theta := (\hat{x}_\theta^2 + \hat{y}_\theta^2)^{1/2}$ . Since the angle  $\theta$  is fixed, this metric represents the aspect of  $(\hat{x}_\theta, \hat{y}_\theta)$  that will vary across different curve estimates. Uncertainty in return curve estimates can consequently be quantified using the distribution of  $\hat{d}_\theta$  at each angle  $\theta \in \Theta$ . We propose the following bootstrap procedure: for  $k = 1, \dots, K$ ,

1. Bootstrap the original data sample to produce a new sample of the same size.
2. For each  $\theta \in \Theta$ , obtain the  $l_2$ -norm for the corresponding point estimate obtained using a given model. Denote this value by  $\hat{d}_{\theta,k}$ .

If temporal dependence is shown to exist in the data, block bootstrapping (Kunsch, 1989) can be used for this procedure. This sampling scheme retains temporal dependence in the resampled datasets, which ensures the additional uncertainty that arises due to the lower effective sample size is accounted for (Politis and Romano, 1994). It is widely used for quantifying uncertainty in environmental applications of extreme value theory; see, for instance, Eastoe and Tawn (2009), Mhalla et al. (2019) and Simpson and Wadsworth (2021).

Given  $\theta \in \Theta$ , we construct empirical estimates of the mean, median, and  $100(1 - \alpha)\%$  confidence intervals for the  $l_2$ -norm values using the sample  $\{\hat{d}_{\theta,k} \mid 1 \leq k \leq K\}$ . Taking  $\alpha = 0.95$ , we estimate the 2.5% and 97.5% quantiles using this sample, which we denote  $\hat{d}_\theta^{0.025}$  and  $\hat{d}_\theta^{0.975}$  respectively. Assuming unbiased estimation,  $\Pr(\hat{d}_\theta^{0.025} \leq d_\theta \leq \hat{d}_\theta^{0.975}) \approx 0.95$ , where

$d_\theta = (x_\theta^2 + y_\theta^2)^{1/2}$  is the  $l_2$ -norm of  $(x_\theta, y_\theta) \in \text{RC}(p) \cap L_\theta$ . Hence, one can show that

$$\Pr \left[ (x_\theta, y_\theta) \in \left\{ (x, y) \in L_\theta \mid d_\theta \in [\hat{d}_\theta^{0.025}, \hat{d}_\theta^{0.975}] \right\} \right] \approx 0.95,$$

implying the set  $\{(x, y) \in L_\theta \mid d_\theta \in [\hat{d}_\theta^{0.025}, \hat{d}_\theta^{0.975}]\}$  defines a confidence region for curve points along the line  $y = \tan(\theta)x$ . Taking the maximum and minimum  $x$  and  $y$  coordinates in this set, we obtain a pointwise confidence region for points along the line segment  $L_\theta$  at each angle  $\theta \in \Theta$ . These pointwise confidence regions, along with the  $x$  and  $y$  coordinates corresponding to the mean and median  $l_2$ -norm values, can be joined together in order of angle to construct estimates that represent mean, median, and 95% confidence interval estimates for the return curve.

Our procedure is illustrated in Figure 5 with  $m = 150$ . The confidence interval width appears to vary over angles in both cases - this is partly explained by implementation of Properties 3.1 - 3.4 in each bootstrap curve estimate. For the inverted logistic copula, the true curve is captured by the estimated confidence region at all angles. For the asymmetric logistic copula, the estimated confidence region only captures the true curve in certain regions of the  $\mathbb{R}_+^2$  plane. This observation indicates some bias may exist for curve estimates from this model. This bias is likely a result of the modelling framework being unable to account for the complex asymmetric structure of this copula at finite levels due to a poor rate of convergence to the limiting angular dependence function introduced in Equation (2.3). Many bivariate extreme value models fail to account for asymmetric dependence structures; see, for instance, Tendijck et al. (2021).

## 4.2 Return curve diagnostic tool

Since the true return curve is unknown in practice, we require a means of evaluating the goodness of fit for a curve estimate,  $\widehat{\text{RC}}(p)$ , obtained from a particular sample. We propose such a technique and illustrate the method using a single data set simulated from a logistic copula (Tawn, 1988) on standard exponential margins. This tool provides a means to assess the accuracy of a




Figure 5: Comparison of median, mean, and 95% confidence interval return curve estimates (green, blue and black dotted lines respectively) to the true return curves (red) for inverted logistic (left) and asymmetric logistic (right) copulas with  $K = 1000$  bootstrap samples.

given curve estimate for a data set with no knowledge of marginal or copula distributions.

Consider the shaded survival regions defined in the left panel Figure 6 for an estimated return curve  $\widehat{RC}(p)$ , where  $p$  is small but  $\widehat{RC}(p)$  is in the range of the data. Regions of the form  $(x, \infty) \times (y, \infty)$  are illustrated at three different points on the curve. The probability of lying within each such region should, by definition, equal  $p$ . To assess this, we consider fixed survival regions for a chosen subset of points on the estimated curve. For convenience, this subset is chosen such that points correspond to the set of angles  $\Theta$ ; again we take  $m = 150$ . This results in the set of points sufficing as a representation of the entire estimated curve, as demonstrated in the right panel of Figure 6. If this estimated curve accurately reflects the true return curve, the empirical probability of observing data within each survival region should be close to  $p$ .

Let  $j \in [1, 150]$  denote the index of an angle  $\theta_j \in \Theta$ , and let  $(\hat{x}_{\theta_j}, \hat{y}_{\theta_j})$  denote the corresponding point on the estimated curve. Furthermore, let  $(\mathbf{x}, \mathbf{y}) = \{(x_i, y_i) : 1 \leq i \leq n\}$  denote the observed sample that has been used to estimate the curve. The empirical estimate, which we denote  $\hat{p}_j$ , is given by the proportion of points lying in the region  $(\hat{x}_{\theta_j}, \infty) \times (\hat{y}_{\theta_j}, \infty)$ . We then apply the bootstrap to resample the original data set and this estimation procedure is repeated to obtain a range of empirical estimates. As in Section 4.1, block bootstrapping should be applied if temporal




Figure 6: Left: Survival regions for three points on the estimated return curve. Right: Subset of points on the estimated curve considered in diagnostic.

dependence is shown to exist in the data set. For each  $j$ , we let  $\hat{\mathcal{P}}_j$  denote the set of empirical probability estimates obtained using bootstrapping. Finally, we estimate the median and 95% pointwise confidence intervals for the probabilities at index  $j$  by taking empirical 2.5%, 50% and 97.5% quantiles of the set  $\hat{\mathcal{P}}_j$ . These estimates provide a pointwise diagnostic at each angle, and can be combined over angles to represent the diagnostic procedure over the whole curve.

This procedure is illustrated in Figure 7 using the example given in Figure 6. The black line and shaded regions in the figure represent the empirical estimates of the median and 95% pointwise confidence intervals, respectively, for each index, with the red line denoting the true probability. As can be observed, for all indices, the confidence bounds contain the true value  $p$ , suggesting this estimated curve accurately represents this value. However, the median empirical estimates are greater than  $p$  at the majority of indices, suggesting a slight overestimation bias for this particular curve estimate.

We note that the confidence intervals produced through bootstrapping for the return curves in Section 4.1, and survival probabilities in Section 4.2, are all pointwise and dependent across angles. Although they cannot be interpreted across the whole range, they still provide a useful assessment of the utility of various curve estimation techniques. Both tools are adapted in Section



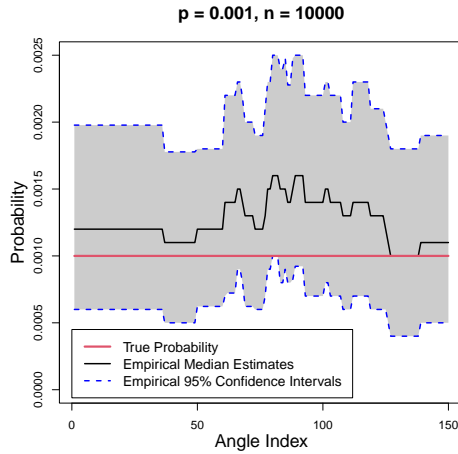


Figure 7: Illustration of diagnostic tool. Solid red and black lines denotes true and mean empirical estimates, respectively, and grey shaded region between dotted blues lines describe empirical 95% CI estimates.

6 to account for the original margins of the environmental data sets, allowing us to analyse the quality of return curve estimates for these examples.

## 5 Simulation study

We compare the return curve estimates from the models discussed in Section 3.4 to those estimated using the methodology of Cooley et al. (2019). For this, we consider several simulated data sets on standard exponential margins, representing a range of different extremal dependence structures. Specifically, we consider the following copula families: logistic and asymmetric logistic copulas from the bivariate extreme value (BEV) family, the bivariate normal copula with correlation coefficient  $\rho$ , logistic and asymmetric logistic copulas from the inverted BEV family, the bivariate t copula with correlation coefficient  $\rho$  and degrees of freedom  $\nu$  and the Frank copula with dependence parameter  $\zeta$ .

For the methods introduced in Cooley et al. (2019), we transform data to standard Fréchet margins, use the procedures proposed in the paper to obtain return curve estimates, transform back to standard exponential margins and apply Properties 3.1, 3.3 and 3.4. For each example,

the chosen estimation procedure is determined by the extremal dependence exhibited by the underlying copula; this must be specified prior to inference, illustrating a drawback of this approach. We use our knowledge of the true dependence structure to implement the correct procedure, but in practice such knowledge would not be available to us. The code for implementing this approach can be found at <https://www.stat.colostate.edu/~cooleyd/Isolines/>.

Examples of both estimated and true return curves for each copula, with  $n = 10000$  and  $p = 10^{-3}$ , are illustrated in Figure 8. We note that for the Frank copula, there is a distinct ‘linear’ segment of the curve estimate from the Heffernan and Tawn (2004) model: this lack of fit reflects a shortcoming of this approach for data sets with negative dependence. Since the modelling framework requires us to condition on either  $X_L$  or  $Y_L$ , we can only evaluate joint tail behaviour in the region where at least one variable is large, i.e.,  $\{(x,y) \in \mathbb{R}^2 \mid x > u_{X_L} \text{ OR } y > u_{Y_L}\}$ . As observed for the Frank copula, part of the true return curve can be defined outside of this region for negatively dependent data sets; consequently, this curve region cannot be estimated using the Heffernan and Tawn (2004) framework. This explains the linear segment, since the point estimates for the regions  $R_{Y_L > X_L}$  and  $R_{Y_L \leq X_L}$  are connected to obtain  $\widehat{RC}(p)$ .

To compare curve estimates, we evaluated bias, computational cost and sampling uncertainty for each of the three curve estimation procedures. To assess bias, 1000 samples of size  $n = 100000$  were simulated from each copula and the probabilities  $p = 10^{-3}$  and  $p = 10^{-4}$  were considered. Adapting the procedure detailed in Section 4.1, median curve estimates for each copula were obtained over the 1000 samples and compared to the corresponding true curves. While a median curve cannot be computed in practice, this measure should provide a summary of the bias that arises from each procedure.

To summarise results, we consider a plot of the  $l_2$ -norm values at angles  $\theta_j \in \Theta$  against the corresponding indices  $j \in [1, m]$  for the true and estimated median curves, with  $m = 150$ . Numerical methods can then be used to compute the absolute area between the resulting norm curves, with smaller area values corresponding to median curves with less bias. An illustration

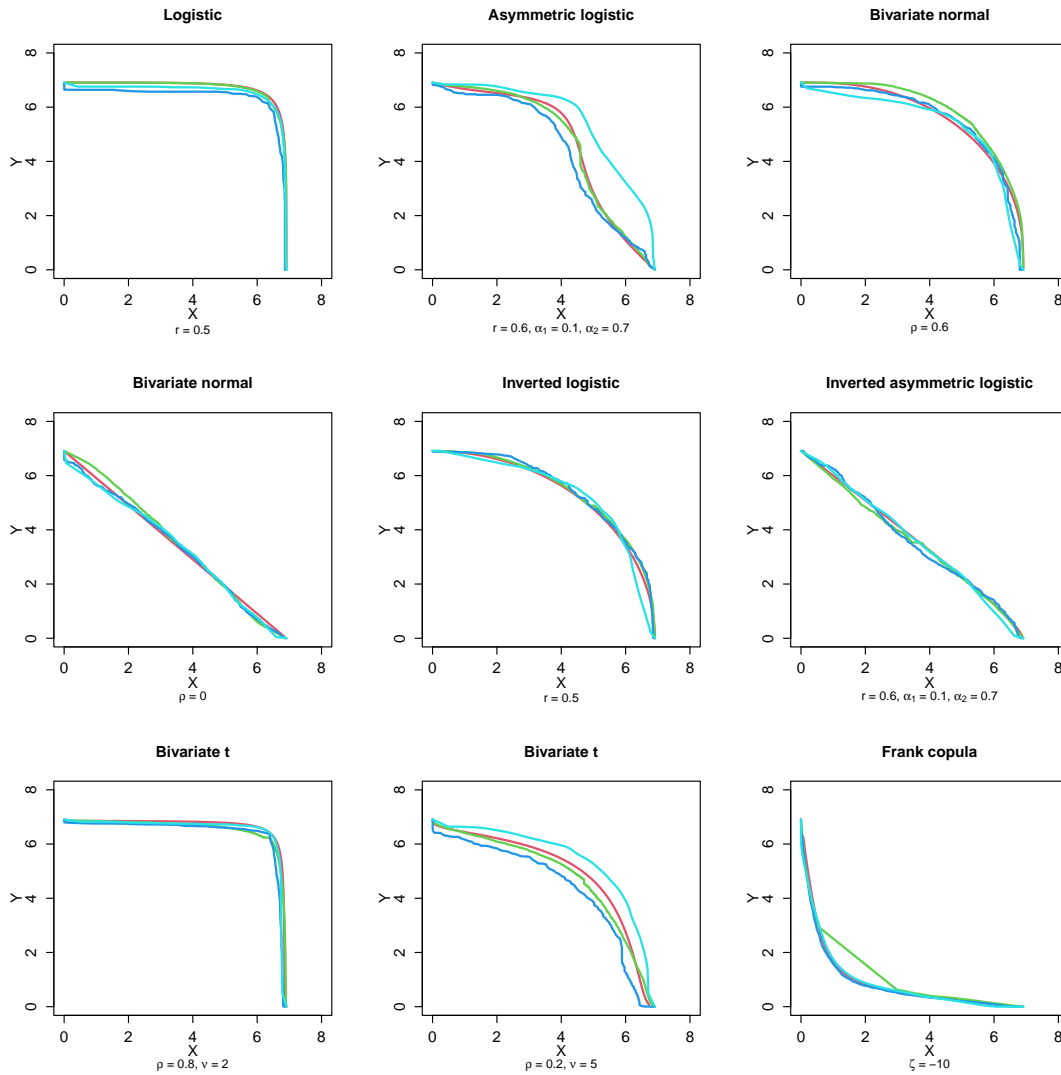


Figure 8: Examples curve estimates for each copula family. True curves are given in red, while the estimated curves from the Heffernan and Tawn (2004), Wadsworth and Tawn (2013) and Cooley et al. (2019) models are given in green, dark blue and light blue, respectively.

of this procedure can be found in the supplementary material.

The summary statistics of integrated absolute difference for each copula-model pairing are given in Table 1. The bias from each procedure appears to vary significantly over the different copula structures, suggesting that the bias in return curve estimates exhibited by a particular model varies with the form of extremal dependence. It is clear that the bias from the Cooley et al. (2019) curve estimates are significantly larger for all but three of the copulas considered; however, we note that in all cases, this method has an unrealistic advantage, namely that the extremal

dependence classes have been correctly specified. On the other hand, the models proposed in Heffernan and Tawn (2004) and Wadsworth and Tawn (2013) appear to have similar amounts of bias across the majority of copula structures considered, and neither consistently outperforms the other. Since models for multivariate extremes are typically based on asymptotic arguments which sometimes hold better for one data set than another, this conclusion is most likely a reflection of the different asymptotic arguments for these models.

To evaluate the computational cost of each estimation technique, fifty samples of size  $n = 10000$  were simulated from a logistic copula with dependence parameter 0.5. With  $p = 10^{-3}$ , a Windows machine with a 1.60 GHz Intel(R) Core(TM) i5-8250U processor and 16GB of RAM was used to compute return curve estimates for each of the fifty samples, and the total computation times were recorded. For the Heffernan and Tawn (2004), Wadsworth and Tawn (2013) and Cooley et al. (2019) techniques, these times were 269.9s, 6.1s, and 2618.2s, respectively.

Application of the Wadsworth and Tawn (2013) modelling framework was significantly quicker than the other two approaches; this is likely due to the fact this technique does not involve any simulation and/or smoothing. Of the remaining estimation frameworks, application of the Heffernan and Tawn (2004) model was still significantly quicker than the method given in Cooley et al. (2019). This conclusion appears to be a result of the Gaussian-kernel density smoothing techniques that are applied when obtaining the empirical curve estimates for the latter approach. Combined with the fact the bias appears significantly lower for the other two models, we choose not consider the approach of Cooley et al. (2019) further.

To assess the sampling uncertainty from the remaining procedures, we compute the coverage for estimated confidence regions at fixed angles. For this, 500 simulated samples of size  $n = 10000$  from each copula were considered. Using bootstrapping with  $K = 200$  iterations, confidence regions were obtained following the procedure outlined in Section 4.1 for probabilities of  $p = 10^{-3}$  and  $10^{-4}$ , and we assessed the coverage of these at five fixed angles  $\theta \in \{\pi(m+1-j)/2(m+1) \mid j = 1, 38, 75, 112, 150\}$ , allowing assessment of coverage for a

variety of regions. Two of the angles are only considered for the BEV asymmetric logistic and inverted BEV asymmetric logistic copulas, since these are the only distributions not to exhibit symmetry. We consider 95% confidence regions for both probabilities. The results for  $p = 10^{-3}$  are given in Table 2; the results for  $p = 10^{-4}$  can be found in the supplementary material, along with an visual illustration of the coverage procedure.

These coverage results provide an insight into differences in the Heffernan and Tawn (2004) and Wadsworth and Tawn (2013) models. Firstly, for angles close to 0 and  $\pi/2$ , the coverage from the Wadsworth and Tawn (2013) model tends to be closer to the nominal level than that from the Heffernan and Tawn (2004) model. This is especially apparent when examining the scores at both probabilities for the logistic and first bivariate normal copula examples. We note that imposing Property 3.1 will affect the coverage near the margins, since we do not allow return curve coordinate estimates that exceed the marginal  $(1 - p)$ -th quantiles, resulting in constrained confidence intervals. We also note that the coverage values for the Frank copula from the Heffernan and Tawn (2004) framework are noticeably small; this relates to aforementioned shortcoming of this approach for data sets with negative dependence. On the other hand, certain coverage values obtained using the Wadsworth and Tawn (2013) approach are noticeably smaller than the corresponding values from the Heffernan and Tawn (2004) approach; for example, for the BEV asymmetric logistic copula at  $p = 10^{-4}$  and the second bivariate t copula. On the whole, neither procedure consistently outperforms the other over the copulas and angles considered and encouragingly, the resulting coverage scores were, in many cases, close to the nominal level.

From these results, we suggest that the curve estimation technique derived using the Wadsworth and Tawn (2013) model is preferable in a practical setting; it is straightforward to implement and significantly outperforms the other considered techniques in terms of computation time. Combined with bias and coverage results, alongside the ability of the Wadsworth and Tawn (2013) model to capture negative dependence structures, this curve estimation technique offers clear advantages over the alternative methods, making it the best suited for practical applications.

## 6 Case study

We now apply our techniques to the two metocean data sets introduced in Section 1. We first transform both data sets to standard exponential margins. Assuming each margin is identically distributed over time, we estimate the marginal distributions using equation (3.1) and apply the probability integral transform. We then use the techniques proposed in Section 3.4 to obtain curve estimates for the probability  $p = 10^{-3}$ , corresponding to a return period of approximately 4.7 years. The resulting curve estimates are illustrated in Figure 9 on the original margins (following back transformation).




Figure 9: Return curves estimated for measured (left) and hindcast (right) data sets, with  $p = 10^{-3}$ . Green and blue lines represent the estimates from the Heffernan and Tawn (2004) and Wadsworth and Tawn (2013) models, respectively.

Next, we adapt the diagnostic tool introduced in Section 4 to assess the goodness of fit of return curve estimates for both data sets. To account for the additional uncertainty that arises during estimation of the marginal distributions, we apply the diagnostic on the original margins of the data. This is done as follows: letting  $\{(x_i, y_i)\}_{i=1, \dots, n}$  denote either data set, we define  $(x_0, y_0) = (\min_i x_i, \min_i y_i)$ . We use these coordinates as a reference point from which we can evaluate return curve estimates. Given  $\Theta$  defined as before, we define the line segment  $L'_\theta := \{(x, y) \in \mathbb{R}^2 \mid y = (x - x_0) \tan(\theta) + y_0\}$  for each  $\theta \in \Theta$  and, for any curve estimate  $\widehat{RC}(p)$ ,

consider the intersection of the sets  $\widehat{RC}(p) \cap L'_\theta$ . Illustrations of reference points, line segments and intersection points are given in the supplementary material for both data sets. Similarly to before, these intersection points are used to define the joint survival regions; the data is then resampled and sets of empirical probability estimates are obtained for each region.

We apply block bootstrapping for resampling, because even with pre-processing, both data sets still appear to exhibit some marginal temporal dependence. Block sizes of 5 and 10 were selected for the measured and hindcast data sets, respectively, by considering plots of the auto-correlation function and selecting values beyond which the dependence appeared insignificant for both variables. These block sizes were then used to bootstrap the original data sets. The resulting diagnostic plots are given in Figures 10 and 11.

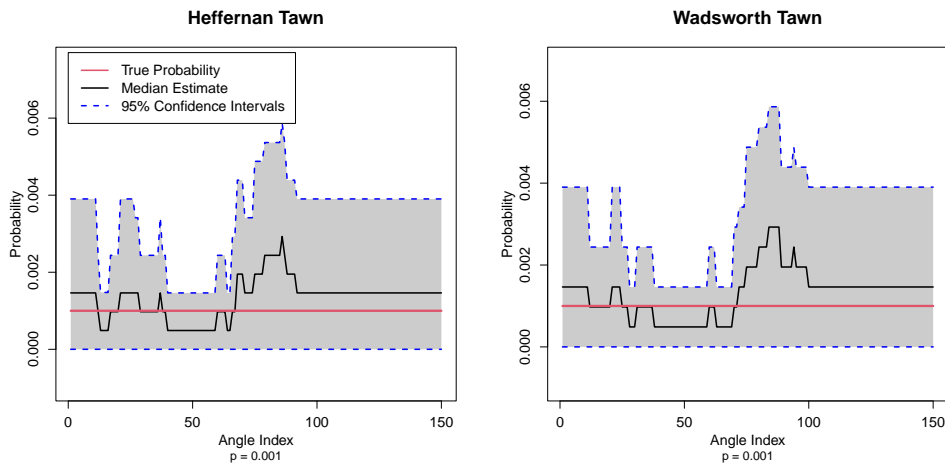


Figure 10: Measured data diagnostic plots with  $K = 1000$  block bootstraps from Heffernan and Tawn (2004) (left) and Wadsworth and Tawn (2013) (right) models, respectively.

For the observed data, Figure 10 suggests both models perform similarly and provide accurate curve estimates, while for the hindcast data, Figure 11 suggests the Wadsworth and Tawn (2013) curve estimate outperforms the Heffernan and Tawn (2004) estimate at a subset of angles. This is reflected by the difference in curve estimates in the the joint upper tail of the data; see Figure 9. In both cases, the variability in curve estimates appears to vary with the angle. The estimated confidence intervals capture the true probability at the majority of considered angles, suggesting

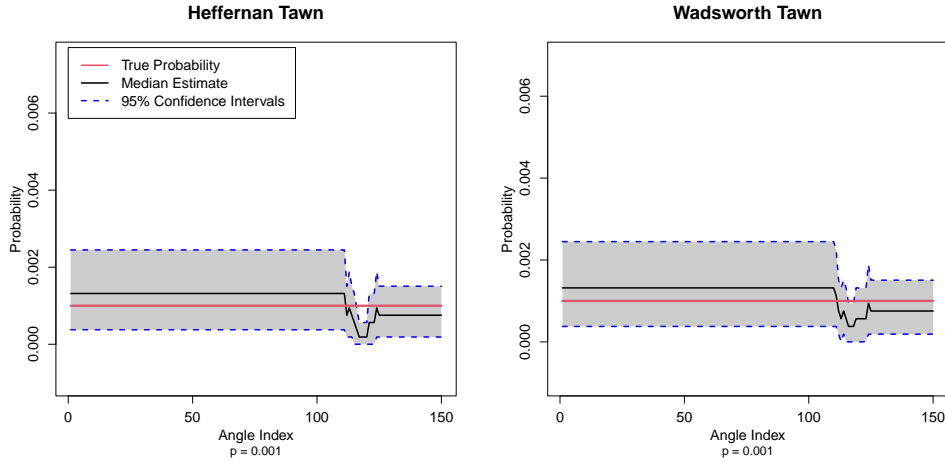


Figure 11: Hindcast data diagnostic plots with  $K = 1000$  block bootstraps from Heffernan and Tawn (2004) (left) and Wadsworth and Tawn (2013) (right) models, respectively.

both estimation techniques can accurately capture joint tail behaviour for these data sets.

Finally, we apply an adaptation of the technique introduced in Section 4.1 to assess uncertainty in return curve estimates. Letting  $(\hat{x}_\theta, \hat{y}_\theta) := \widehat{\text{RC}}(p) \cap L'_\theta$ , with  $L'_\theta$  defined as before, uncertainty is captured by considering the distribution of  $l_2$ -norm values from the reference point, i.e.,  $\hat{d}_\theta = |(\hat{x}_\theta - x_0)^2 + (\hat{y}_\theta - y_0)^2|^{1/2}$ . Using block bootstrapping, with the marginal distributions re-estimated for each bootstrapped sample, we obtain median and mean curve estimates, along with pointwise 95% confidence intervals across angles. For the reasons outlined in Section 4.2, we use the Wadsworth and Tawn (2013) approach to obtain these estimates. The resulting curve estimates are illustrated in Figure 12; we observe that taking mean and median curves appears to have a smoothing effect on the resulting estimates, relative to the original curve estimates given in Figure 9.

## 7 Discussion

We have considered the concept of a return curve as a bivariate extension to a return level and introduced novel estimation techniques, illustrating that these methods perform better than an existing approach. Furthermore, unlike Cooley et al. (2019), our methods do not require the form of



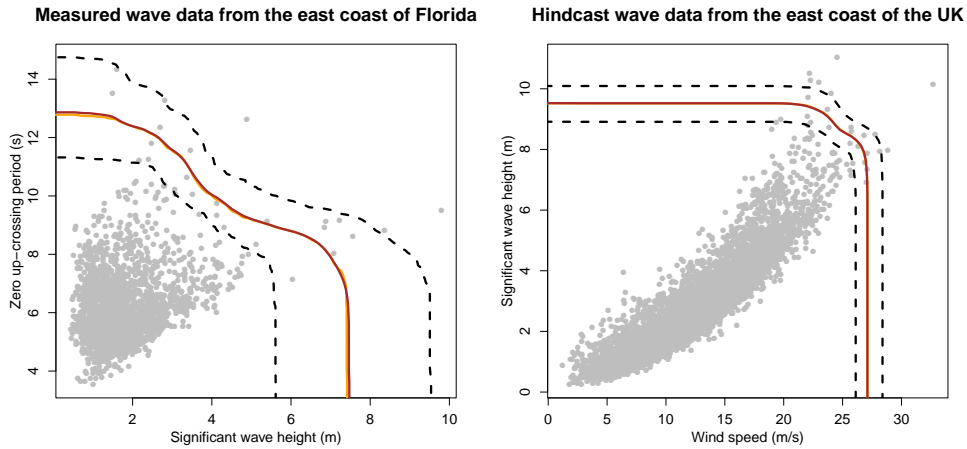


Figure 12: Median (orange) and mean (brown) curve estimate, along with 95% (black dotted) confidence regions obtained using block bootstrapping with  $K = 250$  for the measured (left) and hindcast (right) data sets. The Wadsworth and Tawn (2013) model was used to obtain all estimates.

extremal dependence to be pre-specified: this is an obvious advantage, since determining the extremal dependence structure is seldom straightforward. We have also proposed novel uncertainty representation and diagnostic tools for the risk measure.

For the diagnostic tool proposed in Section 4.2, we note that extreme survival region probabilities are estimated empirically, meaning the accuracy of such estimates will be directly related to the sample size  $n$  and the probability  $p$ . This represents a broader problem within the extremes literature, since, by definition, we will have observed very few extremes values that can be used to verify and justify a given approach. This issue is illustrated further in the supplementary material, where we consider decreasing probabilities for a fixed sample size. However, as with similar analyses, if the tool appears to illustrate a good fit to the data at less extreme probabilities, we can be more confident in extrapolating to more extreme values.

While we have focused on bivariate random vectors, both concepts and methodology can be extended to the general multivariate setting. However, higher dimensional curves are difficult to visualise and capturing the dependence structures in higher dimensions becomes increasingly complex, since different two dimensional marginals can exhibit different forms of extremal de-

pendence. Nevertheless, in the bivariate setting, we believe return curves are a useful tool for researchers to explore joint extremal behaviour and develop a better understanding of potential risks. Indeed, return curves are already utilised to analyse risks for ocean and coastal structures. It is also important to note this risk measure only denotes the rarity of events, not impact: therefore, in practice, researchers must carefully consider which regions of the multivariate space are impactful prior to inference.

As is common in many environmental contexts, the data sets considered in Section 6 both appear to exhibit non-negligible temporal dependence. We account for this feature by using block bootstrapping to quantify uncertainty, but this creates the additional challenge of block size selection. For this, we use an ad-hoc technique based on examinations of ACF plots. An in-depth investigation could improve on this approach through a more robust, theoretically justified resampling scheme.

For the hindcast data set discussed in Section 6, Figure 11 illustrates a downside of applying the diagnostic tools on the original marginal distributions. In particular, the majority of angles selected to represent the curve estimates correspond to the ‘marginal limits’ of the curves, i.e., the straight line segments connecting the curve estimates to the margins in Figure 12. Due to the strong positive dependence between the hindcast variables, empirical survival probabilities will be unchanging along these line segments, explaining the largely constant diagnostic probabilities in Figure 11. This is further demonstrated by the illustration of line segments given in the supplementary material. Future research could explore techniques for selecting angles such that the corresponding return curve representation is exclusively in-between the aforementioned marginal limits.

Finally, we note that all techniques discussed in this paper are only applicable to data sets exhibiting stationarity; accounting for non-stationarity in the context of return curves presents many challenges, since return curves are defined in the stationary setting only and all models introduced in Section 3 assume stationarity. While a range of approaches exist for capturing

non-stationary in the univariate setting (e.g. Eastoe, 2019), relatively few approaches exist in the multivariate setting. This topic has recently been explored in Murphy-Barltrop and Wadsworth (2022), whereby the authors extend the definition of return curves to the non-stationary setting and provide techniques for their estimation.

## Acknowledgements

This paper is based on work completed while Callum Murphy-Barltrop was part of the EPSRC funded STOR-i centre for doctoral training (EP/L015692/1). We are grateful to the two referees for constructive comments that have improved this article.

## SUPPLEMENTARY MATERIAL

### Supplementary material for “New estimation methods for extremal bivariate return curves”

File containing figures and tables that further illustrate ideas and results discussed in the article. (.pdf file)

**Code and data.** Zip file containing two R scripts and the case study data sets. The first script reproduces examples of results from Section 5, while the second script can be used to reproduce case study results from Section 6. (.zip file)

## References

- Balkema, A. A. and de Haan, L. (1974). Residual Life Time at Great Age. *The Annals of Probability*, 2(5).
- Baran, S., Szokol, P., and Szabó, M. (2021). Truncated generalized extreme value distribution-based ensemble model output statistics model for calibration of wind speed ensemble forecasts. *Environmetrics*, 32.

- Coles, S. (2001). *An Introduction to Statistical Modeling of Extreme Values*. Springer Series in Statistics. Springer London, London.
- Coles, S., Heffernan, J., and Tawn, J. (1999). Dependence measures for multivariate extremes. *Extremes*, 2(4):339–365.
- Coles, S. and Pauli, F. (2002). Models and inference for uncertainty in extremal dependence. *Biometrika*, 89(1):183–196.
- Coles, S. G. and Tawn, J. A. (1991). Modelling Extreme Multivariate Events. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 53(2):377–392.
- Cooley, D., Thibaud, E., Castillo, F., and Wehner, M. F. (2019). A nonparametric method for producing isolines of bivariate exceedance probabilities. *Extremes*, 22(3):373–390.
- de Haan, L. and de Ronde, J. (1998). Sea and Wind: Multivariate Extremes at Work. *Extremes 1998 1:1*, 1(1):7–45.
- de Haan, L. and Resnick, S. I. (1977). Limit theory for multivariate sample extremes. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 40(4):317–337.
- Eastoe, E. F. (2019). Nonstationarity in peaks-over-threshold river flows: A regional random effects model. *Environmetrics*, 30(5):1–18.
- Eastoe, E. F. and Tawn, J. A. (2009). Modelling non-stationary extremes with application to surface level ozone. *Journal of the Royal Statistical Society. Series C: Applied Statistics*, 58(1):25–45.
- Eckert-Gallup, A. and Martin, N. (2016). Kernel density estimation (KDE) with adaptive bandwidth selection for environmental contours of extreme sea states. *OCEANS 2016 MTS/IEEE Monterey, OCE 2016*, pages 1–5.

- Environmental Agency (2005). CSG 15 Final Project Report. Technical report, Department for Environment, Food and Rural Affairs.
- Gouldby, B., Wyncoll, D., Panzeri, M., Franklin, M., Hunt, T., Hames, D., Tozer, N., Hawkes, P., Dornbusch, U., and Pullen, T. (2017). Multivariate extreme value modelling of sea conditions around the coast of England. *Proceedings of the Institution of Civil Engineers: Maritime Engineering*, 170(1):3–20.
- Groll, N. and Weisse, R. (2017). A multi-decadal wind-wave hindcast for the North Sea 1949–2014: CoastDat2. *Earth System Science Data*, 9(2):955–968.
- Haselsteiner, A. F., Coe, R. G., Manuel, L., Chai, W., Leira, B., Clarindo, G., Guedes Soares, C., Hannesdóttir, Á., Dimitrov, N., Sander, A., Ohlendorf, J. H., Thoben, K. D., de Hauteclocque, G., Mackay, E., Jonathan, P., Qiao, C., Myers, A., Rode, A., Hildebrandt, A., Schmidt, B., Vanem, E., and Huseby, A. B. (2021). A benchmarking exercise for environmental contours. *Ocean Engineering*, 236:109504.
- Haselsteiner, A. F., Nguyen, P. T., Coe, R. G., Martin, N., Manuel, L., and Eckert-Gallup, A. (2019). A benchmarking exercise on estimating extreme environmental conditions: Methodology and Baseline results. *Proceedings of the International Conference on Offshore Mechanics and Arctic Engineering - OMAE*, 3:1–10.
- Heffernan, J. E. and Tawn, J. A. (2004). A conditional approach for multivariate extreme values. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 66(3):497–546.
- Hill, B. M. (1975). A Simple General Approach to Inference About the Tail of a Distribution. *The Annals of Statistics*, 3(5):1163–1174.
- Huser, R. and Wadsworth, J. L. (2019). Modeling Spatial Processes with Unknown Extremal Dependence Class. *Journal of the American Statistical Association*, 114(525):434–444.

- Jonathan, P., Ewans, K., and Flynn, J. (2014). On the estimation of ocean engineering design contours. *Journal of Offshore Mechanics and Arctic Engineering*, 136(4):1–8.
- Keef, C., Papastathopoulos, I., and Tawn, J. A. (2013). Estimation of the conditional distribution of a multivariate variable given that one of its components is large: Additional constraints for the Heffernan and Tawn model. *Journal of Multivariate Analysis*, 115:396–404.
- Kunsch, H. R. (1989). The jackknife and the bootstrap for general stationary observations. <https://doi.org/10.1214/aos/1176347265>, 17:1217–1241.
- Ledford, A. W. and Tawn, J. A. (1996). Statistics for near independence in multivariate extreme values. *Biometrika*, 83(1):169–187.
- Ledford, A. W. and Tawn, J. A. (1997). Modelling dependence within joint tail regions. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 59(2):475–499.
- Mackay, E. and Jonathan, P. (2020). Assessment of return value estimates from stationary and non-stationary extreme value models. *Ocean Engineering*, 207(April):107406.
- Manuel, L., Nguyen, P. T., Canning, J., Coe, R. G., Eckert-Gallup, A. C., and Martin, N. (2018). Alternative approaches to develop environmental contours from metocean data. *Journal of Ocean Engineering and Marine Energy*, 4(4):293–310.
- Marcon, G., Naveau, P., and Padoan, S. (2017). A semi-parametric stochastic generator for bivariate extreme events. *Stat*, 6(1):184–201.
- Mattei, J., Vial, E., Rebour, V., Liemersdorf, H., and Turschmann, M. (2001). Generic results and conclusions of re-evaluating the flooding protection in French and German nuclear power plants. *Eurosafe*, 1999.
- Mhalla, L., Opitz, T., and Chavez-Demoulin, V. (2019). Exceedance-based nonlinear regression of tail dependence. *Extremes*, 22(3):523–552.

- Murphy-Barltrop, C. J. R. and Wadsworth, J. L. (2022). Modelling non-stationarity in asymptotically independent extremes. *arXiv*, 2203.05860.
- Pickands, J. (1975). Statistical Inference Using Extreme Order Statistics. *The Annals of Statistics*, 3(1):119–131.
- Politis, D. N. and Romano, J. P. (1994). The stationary bootstrap. *Journal of the American Statistical Association*, 89:1303–1313.
- Ramos, A. and Ledford, A. (2009). A new class of models for bivariate joint tails. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 71(1):219–241.
- Resnick, S. (2002). Hidden Regular Variation, Second Order Regular Variation and Asymptotic Independence. *Extremes*, 5(4):303–336.
- Resnick, S. I. (1987). *Extreme Values, Regular Variation and Point Processes*. Springer Series in Operations Research and Financial Engineering. Springer New York, New York.
- Salvadori, G. and De Michele, C. (2004). Frequency analysis via copulas: Theoretical aspects and applications to hydrological events. *Water Resources Research*, 40(12):1–17.
- Serinaldi, F. (2015). Dismissing return periods! *Stochastic Environmental Research and Risk Assessment*, 29(4):1179–1189.
- Shooter, R., Ross, E., Ribal, A., Young, I. R., and Jonathan, P. (2021). Spatial dependence of extreme seas in the north east atlantic from satellite altimeter measurements. *Environmetrics*, 32.
- Simpson, E. S. and Wadsworth, J. L. (2017). Introduction to Extreme Value Theory and Constructing Hazard Curves. Technical report, <https://www.onr.org.uk/documents/2017/onr-rrr-054.pdf>.

- Simpson, E. S. and Wadsworth, J. L. (2021). Conditional modelling of spatio-temporal extremes for red sea surface temperatures. *Spatial Statistics*, 41.
- Stephenson, A. (2002). evd: Extreme Value Distributions. *R News*, 2(2).
- Tawn, J. A. (1988). Bivariate extreme value theory: Models and estimation. *Biometrika*, 75(3):397–415.
- Tendijck, S., Eastoe, E., Tawn, J., Randell, D., and Jonathan, P. (2021). Modeling the Extremes of Bivariate Mixture Distributions With Application to Oceanographic Data.
- Tilloy, A., Malamud, B. D., Winter, H., and Joly-Laugel, A. (2020). Evaluating the efficacy of bivariate extreme modelling approaches for multi-hazard scenarios. *Natural Hazards and Earth System Sciences*, 20:2091–2117.
- Vanem, E., Guo, B., Ross, E., and Jonathan, P. (2020). Comparing different contour methods with response-based methods for extreme ship response analysis. *Marine Structures*, 69(August 2019):102680.
- Velarde, J., Vanem, E., Kramhøft, C., and Sørensen, J. D. (2019). Probabilistic analysis of offshore wind turbines under extreme resonant response: Application of environmental contour method. *Applied Ocean Research*, 93:101947.
- Wadsworth, J. L. and Tawn, J. A. (2013). A new representation for multivariate tail probabilities. *Bernoulli*, 19(5 B):2689–2714.
- Wadsworth, J. L., Tawn, J. A., Davison, A. C., and Elton, D. M. (2017). Modelling across extremal dependence classes. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 79(1):149–175.
- Worms, J. and Naveau, P. (2022). Record events attribution in climate studies. *Environmetrics*.



Yu, C., Blaha, O., Kane, M., Wei, W., Esserman, D., and Zelterman, D. (2022). Regression methods for the appearances of extremes in climate data. *Environmetrics*, 33.

Table 1: Summary statistics for each model under different copula structures. In each case, 1000 samples of  $n = 100000$  datapoints were simulated and the median curves were computed for  $p = 10^{-3}$  and  $10^{-4}$ . ‘HT’, ‘WT’ and ‘CO’ correspond to the median curve estimates from the Heffernan and Tawn (2004), Wadsworth and Tawn (2013) and Cooley et al. (2019) models, respectively.

Copula	Model	$p = 10^{-3}$			$p = 10^{-4}$		
		HT	WT	CO	HT	WT	CO
BEV Logistic		1.80	10.83	0.21	2.56	14.07	0.22
BEV Asymmetric Logistic		23.61	28.27	85.50	47.66	78.16	126.25
Bivariate Normal 1		2.49	3.01	25.31	4.07	7.69	37.21
Bivariate Normal 2		0.18	0.14	8.83	0.37	0.22	13.14
Inverted BEV Logistic		2.75	0.39	26.23	4.06	0.56	34.12
Inverted BEV Asymmetric Logistic		0.80	0.19	11.47	1.22	0.30	11.35
Bivariate T 1		8.26	7.17	2.36	11.24	10.44	3.87
Bivariate T 2		14.62	26.37	65.05	39.77	74.95	90.72
Frank		46.39	9.65	5.65	30.33	46.56	29.70

Table 2: Coverage values of 95% confidence regions for  $p = 10^{-3}$ . ‘HT’ and ‘WT’ correspond to the Heffernan and Tawn (2004) and Wadsworth and Tawn (2013) models, respectively.

Copula	Probability	$p = 10^{-3}$									
	Model	HT					WT				
BEV Logistic	Angle	1	2	3	4	5	1	2	3	4	5
	Coverage	0.014	0.970	0.800	*	*	0.856	0.606	0.896	*	*
BEV Asymmetric Logistic	Angle	1	2	3	4	5	1	2	3	4	5
	Coverage	0.756	0.938	0.868	0.860	0.896	0.908	0.054	0.362	0.902	0.936
Bivariate Normal 1	Angle	1	2	3	4	5	1	2	3	4	5
	Coverage	0.358	0.942	0.956	*	*	0.872	0.930	0.912	*	*
Bivariate Normal 2	Angle	1	2	3	4	5	1	2	3	4	5
	Coverage	0.948	0.930	0.955	*	*	0.930	0.934	0.934	*	*
Inverted BEV Logistic	Angle	1	2	3	4	5	1	2	3	4	5
	Coverage	0.780	0.932	0.932	*	*	0.864	0.944	0.942	*	*
Inverted BEV Asymmetric Logistic	Angle	1	2	3	4	5	1	2	3	4	5
	Coverage	0.960	0.929	0.942	0.951	0.922	0.924	0.934	0.958	0.942	0.904
Bivariate T 1	Angle	1	2	3	4	5	1	2	3	4	5
	Coverage	0.504	0.928	0.512	*	*	0.920	0.758	0.884	*	*
Bivariate T 2	Angle	1	2	3	4	5	1	2	3	4	5
	Coverage	0.784	0.896	0.930	*	*	0.938	0.806	0.632	*	*
Frank	Angle	1	2	3	4	5	1	2	3	4	5
	Coverage	0.896	0	0	*	*	0.922	0.692	0.714	*	*