

Joint modelling of wave spectral parameters for extreme sea states

Philip Jonathan

Shell Technology Centre Thornton
P.O. Box 1
Chester
United Kingdom
philip.jonathan@shell.com

Jan Flynn

Shell International Exploration and Production
P.O. Box 60
2280 AB Rijswijk
The Netherlands
jan.flynn@shell.com

Kevin Ewans

Shell International Exploration and Production
P.O. Box 60
2280 AB Rijswijk
The Netherlands
kevin.ewans@shell.com

September 2009

Abstract

Characterising the dependence between extreme values of wave spectral parameters such as significant wave height (H_S) and spectral peak period (T_P) is of considerable importance in understanding extreme ocean environments and in the design and assessment of marine structures in particular. For example, it is known that mean values of wave periods tend to increase with increasing storm intensity. Here we seek to characterise joint dependence in a straight forward manner, accessible to the ocean engineering community, using a statistically sound approach.

Many methods of multivariate extreme value analyses are based on models which assume implicitly that in some joint tail region each parameter is either independent of or asymptotically dependent on other parameters; yet in reality the dependence structure in general is neither of these. The underpinning assumption of multivariate regular variation restricts these methods to estimation of joint regions in which all parameters are extreme; but regions where only a subset of parameters are extreme can be equally important for design. The conditional approach of Heffernan and Tawn [2004], similar in spirit to that of Haver [1985] but with a more solid theoretical foundation, overcomes these difficulties.

We use the conditional approach to characterise the dependence structure of H_S and T_P . The key elements of the modelling procedure are: (1) marginal modelling for all parameters, (2) transformation of data to a common standard Gumbel marginal form, (3) modelling dependence between data for extremes of pairs of parameters using a form of regression, (4) simulation of long return periods to estimate joint extremes. We demonstrate the approach in application to measured and hindcast data from the Northern North Sea, the Gulf of Mexico and the North West Shelf of Australia. We also illustrate the use of data re-sampling techniques such as bootstrapping to estimate the uncertainty in marginal and dependence models and accommodate this uncertainty in extreme quantile estimation.

We discuss the current approach in the context of other approaches to multivariate extreme value estimation popular in the ocean engineering community.

1 Introduction

Careful statistical description is central to our understanding of extreme ocean environments and reliable design and assessment of ocean structures. Wave climate can be described in terms of sea-state variables such as the significant wave height, H_S , and spectral peak wave period, T_P , or mean zero up-crossing wave period, T_Z , for example. These and other variables together characterise statistical properties of wave climate fully. Spectra for an extreme sea state, such as the 100-year return-period sea state are often required for assessing dynamic loads. We estimate spectral properties of extreme sea states by jointly estimating values of H_S and T_P corresponding to a given return period. If these variables were independent (or effectively so), this procedure would be relatively straight forward - but this is usually not the case. Therefore, reliable estimation of extreme wave climate conditions requires that dependence between different variables be characterised adequately, using some form of multivariate extreme value analysis.

Suppose we wish to model extremes of a sample of values drawn from some multivariate distribution. To achieve this we need to characterise the tail behaviour of the multivariate distribution from which the sample is drawn based on the sample alone. To simplify presentation of mathematical formulae in what follows, we consider modelling extremes of a bivariate random variable (X, Y) (corresponding to (H_S, T_P) , e.g.), but note that ideas discussed are applicable in higher dimensions. Assuming that marginal distributions of individual variables have been transformed to Gumbel scale (by means of the probability integral transform), a key property of the tail of the distribution is its coefficient of tail dependence, η , a constant $\in [0, 1]$, which quantifies the extent of asymptotic dependence for the distribution. For the bivariate random variable (X, Y) , we have:

$$Pr((X, Y) \in A + d) = \exp(-d/\eta)Pr((X, Y) \in A) \quad (1)$$

where $A + d$ is a translation of set A component-wise. This expression is an approximation in general, equivalent to assuming multivariate regular variation in Gumbel marginals. The value of η is fixed for a given distribution, sometimes a function of the parameters of that distribution. When $\eta=1$ we say that the bivariate distribution is asymptotically dependent. Otherwise, it is asymptotically independent. A large number of distributions appropriate for modelling extremes of multivariate samples has been proposed, each with its own parameters and value for η in particular. Modelling our sample therefore requires (1) fitting our distribution of choice, (2) estimating its parameters and η , and (3) confirming goodness of model fit, and various modelling assumptions made, using diagnostic tests.

It has been shown by numerous authors (e.g. Reiss and Thomas 2001), Ledford and Tawn [1996], Ledford and Tawn [2003]) that this approach is adequate for many applications in low dimensions. The approach is valid when the distributional form is chosen appropriately. However, it is often difficult to decide between one of a number of different parametric forms (with possibly different asymptotic dependence structures). Some authors, especially in the engineering literature, neglect to estimate the coefficient of tail dependence.

Ferreira and Guedes-Soares [2002] present a method to estimate bivariate distributions of significant wave height, H_S and mean wave period, T_Z . The method consists of (1) transforming individual variables using the Box-Cox transformation so that, marginally, transformed values are consistent with samples from a normal distribution, (2) fitting a bivariate Normal model to transformed data. The approach is illustrated in application to modelling of Waverider buoy data at Figuera da Foz, Portugal. The approach will be appropriate for modelling the body of a bivariate distribution, provided that Box-Cox transformations provide favourable results. However, the method imposes a particular structure to *extreme tail behaviour* which may not be supported by the sample. In a similar vein, methods of (multivariate) kernel density estimation, appropriate for modelling the body of a distribution, impose the kernel structure on extreme tails also with little or no justification.

Even when carefully applied, there are short-comings to direct fitting of multivariate distributions for extreme values. Firstly, the approach is not suitable to model samples which are not simultaneously extreme in all their components. This is clear since, in such situations, the set $(X, Y) \in A + d$ (in Equation 1) is likely to be empty if $(X, Y) \in A$ is not. Secondly, the approach becomes unmanageable when dimensionality of data increases. For example, we might wish to jointly model extremes of wave spectral parameters (H_S, T_P) over 1000 locations, requiring a 2000-dimensional model, which would not easily be estimated without making simplifying assumptions. Other authors have taken a different, conditional, approach, the basis of which is to take each variable in turn as a conditioning variate, and model remaining variables conditional on a

large value of the conditioning variate. In our bivariate case we might model $Y|X = x$ for large x , and $X|Y = y$ for large y . The advantage of this approach is that it extends in principle to high dimensions, and to applications to samples for which not all variables are simultaneously extreme. A disadvantage of the approach, until recently, was that the form of the conditional model had to be assumed (or itself estimated) from the sample. For example, Haver [1985] proposed a model for extremes of (H_S, T_P) data, based on (1) modelling the extreme tail of H_S using a Weibull form, (2) modelling the distribution of T_P , for a given (range of) H_S using a log-normal form, (3) estimating values for the parameters in the log-normal distribution for $T_P|H_S = h$ corresponding to extreme values h (beyond the sample) by regressing the log-normal parameter estimates on H_S for the sample and extrapolating, and hence estimating extreme quantiles of (H_S, T_P) . Justification for the choice of log-normal conditional model, and for the forms of models for log-normal parameters as a function of H_S was goodness of fit to the sample; clearly, this does not guarantee that models will be adequate when extrapolated beyond the sample.

The conditional approach of Heffernan and Tawn [2004] goes some way to overcoming these difficulties. Based on asymptotic arguments, they derive a parametric equation for the form for one variable conditional on a large value of another, valid for extremes from a wide class of multivariate distributions with Gumbel marginals. That is, they provide a solid rationale for a particular choice of parametric form for conditional modelling. The Heffernan and Tawn conditional approach consists of (1) transforming each variable so that marginally its values are consistent with a sample from the standard Gumbel distribution, (2) estimating the parameters of the parametric model for values of one variable conditional on large values of the other, (3) simulating (and back-transforming) to characterise extremal behaviour of the joint distribution (on the original scale). To emphasise, the major advantage of the Heffernan and Tawn conditional approach (henceforth referred to as the "conditional approach") over previous conditional models, is that the functional forms for marginal fitting and conditional modelling have been justified theoretically using asymptotic arguments.

The paper is structured as follows. In Section 2 we introduce four applications motivating current work. In Section 3 we present an overview of the conditional model, and evaluate its performance on simulated examples in Section 4. We turn to analysis of the four applications in Section 5. Conclusions and recommendations are given in Section 6.




Figure 1: Storm peak values of T_P versus H_S for measured NNS storm peak data.

2 Motivating applications

We motivate current work by considering applications to estimation of extreme values of H_S and T_P at a location in the Northern North Sea (NNS) for which both measured and hindcast data are available, to buoy data from a location in the Gulf of Mexico (GoM), and to hindcast data on the North West Shelf (NWS) of Australia. Data correspond to storm peak values for H_S over threshold, observed during periods of storm events, and corresponding values for T_P throughout unless stated otherwise. For the measured NNS example, illustrated in Figure 1, the sample corresponds to 620 storm peak pairs for the period (March 1973, December 2006) measured using a laser device. Hindcast data for the same NNS location (see Figure 2) corresponds to 827 pairs of values for the period (November 1964, April 1998). GoM data in Figure 3 are National Data Buoy Center measurements from buoy 42002 corresponding to 505 pairs for the period (January 1980, December 2007). Finally, NWS hindcast data (shown in Figure 4) correspond to 145 pairs of storm peak values for the period (February 1970, April 2006). All samples exhibit positive dependence between H_S and T_P .

The four application data sets are chosen to illustrate use of the conditional model, not for development of design criteria. In general, sample sizes are only of the order of 30 years in length, and data include strong covariate effects due, for example, to seasonality. No attempt has been made at this point to accommodate these effects in the analysis. Accordingly, the results shown in section 5 are not intended to be valid estimates of design criteria for those locations.

Distributions of threshold exceedences of storm peak H_S and T_P are consistent with random samples from Generalised Pareto (GP) distribution (see Section 3 below). For example, goodness of fit of the GP distribution, to values of storm peak H_S and T_P respectively for the measured NNS sample is illustrated as a quantile-quantile plot in Figures 5 and 6. The marginal extremal behaviour of storm peak T_P is of course not generally of interest for engineering purposes, but it is important to establish its conformance to GP for the present application. Fits to other samples are of comparable quality. In estimating the GP model, it is important to explore the stability of estimated GP shape and scale parameters as a function of threshold; this is illustrated in Figures 7 and 8 for the measured NNS sample. It can be seen that the GP shape and scale parameters are more stable with respect to threshold for H_S than for T_P in this case. A threshold corresponding to a non-exceedence probability in the region of 0.6 was used for marginal modelling for all




Figure 2: Storm peak values of T_P versus H_S for hindcast NNS storm peak data.




Figure 3: Storm peak values of T_P versus H_S for measured GoM storm peak data.




Figure 4: Storm peak values of T_P versus H_S for hindcast NWS storm peak data.

samples, notwithstanding the fact that threshold selection is a non-trivial step in general.

3 Outline of theory

For pairs (X, Y) of random variables with marginal Gumbel distributions, Heffernan and Tawn [2004] derive a single parametric form for the conditional distribution of one variable given a large value of the other. This parametric form, motivated by the assumption of a particular limit representation for the conditional distribution (see Equation 3 in Section 6), is shown using asymptotic expansions, to be appropriate to characterise the conditional behaviour of a wide range of theoretical examples of bivariate (and higher-dimensional) models for extremes. Model form for positively associated pairs of variables (X, Y) takes a particularly simple form:

$$(Y|X = x) = ax + x^b Z \quad (2)$$

where a and b are location and scale parameters, respectively, to be estimated. Z is a standardised variable, independent of X , converging with increasing x to a non-degenerate limiting distribution G . Extremal dependence behaviour is then characterised by a , b and G . The form of distribution G is not specified by theory.

For a sample $\{x_i, y_i\}_{i=1}^n$ of values from (X, Y) with values of X exceeding an appropriate threshold u , the values of a , b and G are estimated using regression. For simplicity and computational ease, G is assumed to be a Gaussian distribution with mean μ_Z and variance σ_Z^2 treated as nuisance parameters in the model. Fitted values:

$$\hat{z}_i = \frac{(y_i - \hat{a}x_i)}{x_i^{\hat{b}}}, i = 1, 2, 3, \dots, n$$

provide an estimate of a sample from distribution G , useful for subsequent simulations. The adequacy of model fit can be assessed, for example, by (1) demonstrating that the values $\{\hat{z}_i\}_{i=1}^n$ and $\{x_i\}_{i=1}^n$ are not obviously dependent (thus violating a modelling assumption), (2) exploring the effect of varying u on \hat{a} , \hat{b}




Figure 5: Quantile-quantile plot for GP fit to measured NNS storm peak H_S with a threshold non-exceedence probability of 0.6 for modelling, and estimated parameters $\hat{\xi}=-0.11$, $\hat{\beta}=1.5$. Only points with non-exceedence probabilities >0.9 shown.




Figure 6: Quantile-quantile plot for GP fit to measured NNS storm peak T_P with a threshold non-exceedence probability of 0.6 for modelling, and estimated parameters $\hat{\xi}=-0.19$, $\hat{\beta}=2.0$. Only points with non-exceedence probabilities >0.9 shown.




Figure 7: Variation of GP shape estimate $\hat{\xi}$ with threshold for measured NNS storm peak H_S (dots) and T_P (circles).




Figure 8: Variation of GP scale estimate $\hat{\beta}$ with threshold for measured NNS storm peak H_S (dots) and T_P (circles).

and subsequent estimates (e.g. of probabilities associated with extreme sets), and (3) bootstrap resampling to estimate the uncertainty of estimates for a , b and subsequent estimates for a given threshold choice.

The Generalised Pareto (GP) form is appropriate for model marginal distributions of peaks over threshold, rather than the Gumbel distribution. Therefore, in such cases, to use the conditional model above we transform original variables (X^*, Y^*) from GP to Gumbel using the probability integral transform as follows. Suppose we fit the GP distribution (to the sample from X^* without loss of generality):

$$F_{GP}(x; \xi, \beta, u) = 1 - (1 + \frac{\xi}{\beta}(x - u))_+^{-\frac{1}{\xi}}$$

and estimate cumulative probabilities $\{F_{GP}(x_i^*; \hat{\xi}, \hat{\beta}, u)\}_{i=1}^n$. Now the standard Gumbel distribution has cumulative distribution function:

$$F_{Gmb}(x) = \exp(-\exp(-x))$$

Thus if we define the transformed sample $\{x_i\}_{i=1}^n$ such that $F_{Gmb}(x_i) = F_{GP}(x_i^*; \hat{\xi}, \hat{\beta}, u)$ or $x_i = -\log(-\log(F_{GP}(x_i^*; \hat{\xi}, \hat{\beta}, u)))$ for $i = 1, 2, 3, \dots, n$, the transformed sample will be consistent with a sample from a Gumbel distribution. Similarly, given a value of x (of Gumbel variate), we can calculate the corresponding value x^* on the original GP scale.

It is essential to demonstrate the adequacy of GP marginal fits, e.g. by using the mean residual life plot and stability of marginal shape parameter as a function of threshold choice.

Finally, to simulate a random drawing from the conditional distribution $Y|X > u$, the following procedure can then be followed: (1) draw a value x of X at random from its standard Gumbel distribution, given that the value exceeds threshold u (2) draw a value of z of Z at random from the set $\{\hat{z}_i\}_{i=1}^n$ (3) calculate that value of $y|x = \hat{a}x + x^{\hat{b}}z$ (4) transform (x, y) to (x^*, y^*) using the probability integral transform and the estimated GP marginal model parameters

Using simulation, estimates for various extremal statistics (e.g. values associated with long return periods) can be obtained routinely.

4 Evaluation of conditional model

A comprehensive demonstration of the performance of the conditional model has been provided in Hefernan and Tawn [2004]. It is demonstrated that the conditional model works well for both simultaneous and non-simultaneous bivariate extremes from various distributions. In particular, the approach naturally accommodates different degrees of asymptotic dependence, quantified in terms of coefficient of tail dependence η . It was shown that the conditional model performs at least as well in general as direct fitting of a multivariate extreme value distribution, especially when the coefficient of tail dependence is not explicitly estimated. Here, we evaluate the performance of the conditional model in application of bivariate modelling of samples from the following bivariate distributions, and compare performance of the conditional model with that of a variant of the approach proposed by Haver (1985).

We perform simulation using the following four distributions, selected since they exhibit different asymptotic dependence structures which we might anticipate observing in oceanographic applications:

Distribution D1. EVEL: Multivariate extreme value distribution with exchangeable logistic dependence

$$Pr(X < x, Y < y) = \exp(-(\exp(-x/\alpha) + \exp(-y/\alpha))^\alpha)$$

where X and Y have standard Gumbel marginals. Conditionally, $Pr(Y|X = x)$ can be written in closed form, so that simulations from this bivariate distribution are straight forward by first simulating one variate from its marginal distribution and subsequently simulating the second variate from the conditional distribution given the value of the first variate. X and Y are dependent (unless $\alpha=1$). It can be shown that $Pr((X, Y) \in A + d) = \exp(-d)Pr((X, Y) \in A)$ regardless of the value of α , i.e. that the coefficient of tail dependence, η , is unity, so that this distribution is by definition asymptotically dependent. It can also be shown, conditionally, that $(Y|X = x) = x + Z$ for sufficiently large x , so this distribution is *asymptotically* conditionally dependent (since $Y|X = x$ depends on x for large x). The value of α has no effect on the

asymptotic conditional dependence structure. The true values of a and b in the conditional form 1 are therefore 1 and 0 respectively.

Distributions D2 and D3. TNrm: Multivariate distribution with Normal dependence transformed marginally to standard Gumbel

$$(X, Y) = -\log(-\log(\Phi_{\Sigma}(W)))$$

where W is a bivariate random variable, $W \sim N(0, \Sigma)$, $\Sigma_{11} = \Sigma_{22} = 1$, $\Sigma_{12} = \Sigma_{21} = \rho$ and Φ_{Σ} is the $N(0, \Sigma)$ cumulant. X and Y have standard Gumbel marginals. X and Y are also dependent (unless $\rho=0$). Conditionally, $Y|X = x \sim N(\rho x, (1 - \rho^2))$, so that simulations from this bivariate distribution are also straight forward. It can further be shown that $Pr(X, Y) \in A + d = \exp(-d/\eta)Pr((X, Y) \in A)$ where $\eta = (1 + \rho)/2$. By definition, this distribution is therefore asymptotically independent (since $\eta \neq 1$ unless $\rho=1$). Conditionally, $(Y|X = x) = \rho^2 x + x^{1/2}Z$ for large x , so the distribution is asymptotically conditionally dependent. The value of ρ affects the conditional dependence structure. The true values of a and b in the conditional form 1 are ρ^2 and 0.5 respectively.

For distribution D2, we set the value of ρ to 0.9 corresponding to strong conditional dependence, and for distribution D3, the value of ρ is set to 0.5.

Distribution D4. AC: Asymptotic conditional form

We define this bivariate distribution directly in terms of the marginal of the conditioning variate X , taken to be standard Gumbel, and the conditional distribution of $Y|X = x$, taken to follow the form $(Y|X = x) = ax + x^b Z$ with $a=0.5$, $b=0.25$ and $Z \sim N(0, 1)$. The true values of a and b in the conditional form 1 are therefore 0.5 and 0.25 respectively. Given that X is Gumbel-distributed, the conditioned variate is only asymptotically Gumbel-distributed. In order that Y also be Gumbel-distributed, we replace its values with a random sample from a standard Gumbel distribution such that rank order of values is preserved.

A key step in the modelling procedure remains estimation of marginal distributions for X and Y . For this reason, we choose to transform simulated samples to generalised Pareto marginals with (shape, scale, threshold) triplets of (-0.15, 2.0, 7.0) and (-0.2, 1.0, 9.0) for X and Y prior to analysis, so that all analysis steps required in practice are also followed in our evaluation procedure. In particular, we will be able to assess the relative importance of estimating marginal and conditional models.

The Haver model (from Haver [1985], referred to therein as a conditional model also, but referred to here as "the Haver model" for clarity), developed explicitly for joint modelling of X (*i.e.* H_S) and Y (*i.e.* T_P) for large values of X , assumes that $Y|X = x$ follows a log-normal distribution:

$$f_{Y|X}(y|x) = \frac{1}{\sqrt{2\pi}\sigma(x)} \frac{1}{x} \exp\left(-\frac{(\log x - \mu(x))^2}{2\sigma^2(x)}\right)$$

Estimated parameters $\mu(x)$ and $\sigma(x)$ from the log-normal fits are then regressed on x , so that extrapolations to larger x (beyond the sample) are possible. Haver recommends functional forms $\mu = \zeta_1(x + 1)^{-2} + \zeta_2 \log_e(x + 1) + \zeta_3$ and $\sigma = \zeta_4 x^{\zeta_5}$, based on their adequacy in fitting data for his applications. No further justification for these functional forms is claimed, and there is no reason to expect that these functional forms will provide adequate extrapolations for large x in general. In the original work, a Weibull model for X over a suitably large threshold, was also estimated. Here, for simplicity and fair comparison, we consider a modified approach in which the marginal model for X is assumed GP. Further, during simulation experiments it was found by inspection of diagnostics that the functional form $\zeta_1 x^{-2} + \zeta_2$ was adequate in general for describing the variation of log-normal parameter estimates within the body of the sample; therefore this form was adopted to describe both $\mu(x)$ and $\sigma(x)$, although no claims are made about the relevance of this form in general, particularly for extrapolation.

To assess performance of the conditional model, and compare it with that of the Haver model, 100 realisations of samples of size 1000 from the four distributions are simulated and modelled. Quality of fit is evaluated in terms of estimates for return value Y_{10} , defined as follows: (1) Consider the most probable return value corresponding to a sample 10-times the size of the actual sample, referred to as X_{10} . We are interested in quantifying the distribution of the value of Y (namely Y_{10}) associated with X_{10} . We summarise the distribution of Y_{10} in terms of estimates for its three quartiles: median, 25th and 75th percentiles. To emphasise the difference between the true and estimated values of these quartiles, we report results in terms of the bias (known value minus estimate) for each quartile. Thus, in Table 1, $q(0.25)$ is the bias of the 25th

(a) Conditional model

Data \ Bias	q(0.25)	q(0.50)	q(0.75)
D1: EVEL	-0.33 (-0.39, -0.24)	-0.23 (-0.27, -0.18)	-0.15 (-0.19, -0.10)
D2: TNrm($\rho=0.9$)	0.04 (-0.06, 0.09)	-0.00 (-0.06, 0.07)	-0.01 (-0.08, 0.05)
D3: TNrm($\rho=0.5$)	0.18 (-0.08, 0.47)	0.13 (-0.18, 0.37)	0.04 (-0.21, 0.24)
D4: Conditional	-0.09 (-0.30, 0.03)	-0.09 (-0.22, 0.03)	-0.08 (-0.18, 0.04)

(b) Haver model

Data \ Bias	q(0.25)	q(0.50)	q(0.75)
D1: EVEL	-1.27 (-1.40, -1.11)	-0.97 (-1.04, -0.85)	-0.60 (-0.67, -0.52)
D2: TNrm($\rho=0.9$)	-0.84 (-0.95, -0.74)	-0.71 (-0.79, -0.63)	-0.53 (-0.62, -0.45)
D3: TNrm($\rho=0.5$)	-0.43 (-0.54, -0.34)	-0.52 (-0.63, -0.39)	-0.52 (-0.69, -0.38)
D4: Conditional	-1.00 (-1.15, -0.85)	-0.92 (-1.07, -0.79)	-0.73 (-0.88, -0.61)

(c) Conditional parameter estimates

Data \ Bias of estimate	a	b
D1: EVEL (a=1.00, b=0.00)	-0.18 (-0.21, -0.16)	0.24 (0.12, 0.38)
D2: TNrm($\rho=0.9$) (a=0.81, b=0.50)	-0.00 (-0.07, 0.06)	-0.07 (-0.18, 0.06)
D3: TNrm($\rho=0.5$) (a=0.25, b=0.50)	0.07 (0.00, 0.17)	-0.29 (-0.42, -0.18)
D4: Conditional (a=0.50, b=0.25)	-0.02 (-0.08, 0.00)	-0.01 (-0.14, 0.08)

Table 1: Estimation when marginal models assumed known: bias in estimates for Y_{10} from (a) the conditional model and (b) the Haver model; (c) estimated bias of parameter estimates for conditional model (see equation 1; true values for parameters given with data description in first column).

percentile of Y_{10} , and $q(0.50)$ is the bias of the median value of Y_{10} . For each of the 100 realisations of data, the estimated bias will also vary. We summarise the distribution of biases $q(0.25)$, $q(0.50)$ and $q(0.75)$ in terms of their median, 25th and 75th percentiles. Thus, for distribution D2 using the conditional model, the median values of $q(0.25)$, $q(0.50)$ and $q(0.75)$ are respectively 0.04, 0.00 and -0.01 to two decimal places. For the Haver model, the corresponding values are -0.84, -0.71 and -0.53. For distribution D2 using the conditional model, bias $q(0.50)$ has an inter-quartile range (-0.06, 0.07). For the conditional model only, we similarly compare (see Table 1(c)) known and estimated values for parameters a (and b). Defining biases $q(0.25)$, $q(0.50)$ and $q(0.75)$ as the difference between the known and estimated value of the 25th percentile, median and 75th percentile of a (and b), distributions of biases are summarised in terms of their median and inter-quartile ranges estimated from 100 sample realisations for each distribution. Thus the median value for bias $q(0.50)$ of a is 0.07 for distribution D3, with an inter-quartile range of (0.00, 0.17). *Note that the marginal distribution of the conditioning variate X is assumed known in this part of the study.* A non-exceedance probability threshold of 0.9 (on the conditioning variate) was used for conditional modelling (that is, retaining only the largest 10% of sample values).

The Haver model was applied as follows. Values of X and Y corresponding to the largest 20% of values of conditioning variate X were sorted by increasing value of X and partitioned into 10 sub-samples of size 20. For each sub-sample, the range of values of X was regarded as being sufficiently small that constant values of log-normal parameters μ and σ would be appropriate. The log-normal model was then applied to each sub-sample in turn yielding 10 estimates for μ and σ as a function of (increasing sub-sample median) value of X , as illustrated in Figures 10 and 11 respectively). The parametric form $\zeta_1 x^{-2} + \zeta_2$ was then applied to model each of the log-normal parameters in turn, and provide estimates for values of Y associated with extreme quantiles of X .

From Table 1(a), we see that in all but case D1, the inter-quartile range for the bias in quantile estimate includes zero, and conclude that the conditional model performs well in general. The variation in parameter estimates \hat{a} and \hat{b} during simulation, summarised in Table 1(c), is also illustrated in Figure 9. Convex hulls enclose sets corresponding to each of the four distributions D1-D4. It can be seen in general that parameter estimates are close to their true values. For distribution D1, both parameter estimates are biased. This lack of fit is reflected also in the relatively large bias in estimation of Y_{10} for this case. Comparing Tables 1(a) and 1(b), we see that the conditional model performs considerably better than the Haver model for the four cases




Figure 9: Variation in parameter estimates \hat{a} and \hat{b} during simulation, *assuming marginals known*, corresponding to results in Table 1. Convex hulls enclose pairs corresponding to distribution D1 (circles), D2(dots), D3(squares) and D4(crosses). Location of true values indicated using numbers 1-4.

examined. There is a systematic bias in the quantile estimates from the Haver model. Of course, by choosing a different form for the conditional model, or a different form for modelling the log-normal parameters as a function of the conditioning variate, this bias could almost certainly be reduced. The models for log-normal mean and standard deviation are illustrated in Figures 10 and 11 for a typical realisation from distribution D1. The essential point here however is that Haver model gives no sound basis for selecting appropriate functional forms for conditional modelling and parameter extrapolation, whereas the current conditional model does.

To investigate the influence of marginal model estimation on overall performance of the conditional model, the simulation study was repeated, *now with both marginal model for X and conditional model for $Y|X$ estimated*. A non-exceedance probability threshold of 0.8 was used for marginal modelling, and a non-exceedance probability threshold of 0.9 for conditional modelling. The goodness of fit is again assessed in terms of the bias of the estimated value of Y_{10} . We also compare known and estimated values for parameter a and b . Median, 25th and 75th percentiles for the sample of bias estimates are given.

Tables 1 and 2 show that estimates for conditional model parameters \hat{a} and \hat{b} are relatively insensitive to estimation of marginal models in general. Unsurprisingly, perhaps, estimation of the index parameter b is more difficult than that of the linear coefficient a when marginal analysis is incorporated. For simulations of large samples (e.g. corresponding to long return periods) using the conditional model (see Section 3), values of residuals of Z from model fitting are sampled as a non-parametric estimate for the (unknown) distribution G for Z . In particular, it is likely that a relatively biased estimate \hat{b} will be compensated by the corresponding sample of residuals of Z during simulation, at least to some extent. Estimated values for \hat{a} and \hat{b} are illustrated in Figure 12.

To assess the overall performance of the model, including both marginal and conditional estimation, but excluding threshold choice for marginal and conditional modelling, we use the bootstrap scheme recommended in Section 5.4 of Heffernan and Tawn [2004]. Briefly, the uncertainty in conditional model parameters \hat{a} and \hat{b} (and others, for example, that of Y_{10}) is estimated by (1) creating new bootstrap samples using the fitted model (both marginal and conditional components), (2) fitting the model to the bootstrap samples




Figure 10: Extrapolation of log-normal mean parameter μ , for a typical realisation of data. Estimated values from the sample (circles joined by solid lines), log-normal fit (dashed line) and estimated extrapolated value for μ (dot)




Figure 11: Extrapolation of log-normal standard deviation parameter σ , for a typical realisation of data. Estimated values from the sample (circles joined by solid lines), log-normal fit (dashed line) and estimated extrapolated value for σ (dot)

(a) Conditional model

Case \ Bias	q(0.25)	q(0.50)	q(0.75)
D1: EVEL	-0.37 (-0.53, -0.22)	-0.32 (-0.47, -0.14)	-0.28 (-0.40, -0.09)
D2: TNrm($\rho=0.9$)	-0.00 (-0.20, 0.19)	-0.04 (-0.25, 0.16)	-0.09 (-0.30, 0.15)
D3: TNrm($\rho=0.5$)	0.14 (-0.08, 0.40)	0.04 (-0.15, 0.29)	-0.01 (-0.24, 0.22)
D4: Conditional	-0.23 (-0.62, 0.03)	-0.19 (-0.46, 0.04)	-0.21 (-0.46, 0.04)

(b) Haver model

Case \ Bias	q(0.25)	q(0.50)	q(0.75)
D1: EVEL	-1.24 (-1.38, -1.12)	-0.96 (-1.04, -0.85)	-0.60 (-0.71, -0.49)
D2: TNrm($\rho=0.9$)	-0.86 (-0.97, -0.75)	-0.73 (-0.82, -0.64)	-0.55 (-0.65, -0.45)
D3: TNrm($\rho=0.5$)	-0.44 (-0.57, -0.34)	-0.53 (-0.63, -0.40)	-0.53 (-0.68, -0.39)
D4: Conditional	-1.01 (-1.12, -0.88)	-0.95 (-1.04, -0.81)	-0.73 (-0.84, -0.60)

(c) Conditional parameter estimates

Case \ Bias of estimate	a	b
D1: EVEL ($a=1.00, b=0.00$)	-0.18 (-0.20, -0.17)	0.25 (0.14, 0.35)
D2: TNrm($\rho=0.9$) ($a=0.81, b=0.50$)	-0.01 (-0.08, 0.06)	-0.07 (-0.18, 0.05)
D3: TNrm($\rho=0.5$) ($a=0.25, b=0.50$)	0.06 (0.00, 0.14)	-0.27 (-0.40, -0.16)
D4: Conditional ($a=0.50, b=0.25$)	-0.05 (-0.25, 0.17)	0.15 (-0.07, 0.34)

Table 2: Estimation when marginal models are estimated: bias in extreme quantile estimates for Y_{10} from (a) the conditional model and (b) the Haver model; (c) estimated bias of parameter estimates for conditional model (see equation 1; true values for parameters given with data description in first column).




Figure 12: Variation in parameter estimates \hat{a} and \hat{b} during simulation, *with marginal estimation*, corresponding to results in Table 2. Convex hulls enclose pairs corresponding to distribution D1 (circles), D2(dots), D3(squares) and D4(crosses). Location of true values indicated using numbers 1-4.




Figure 13: Bootstrap convex hulls for conditional parameters a and b corresponding to typical realisations from distributions D1-D4. Convex hulls enclose pairs corresponding to distribution D1 (circles), D2(dots), D3(squares) and D4(crosses). Location of true values indicated using numbers 1-4.

and estimating parameters \hat{a} and \hat{b} , then (3) quantifying the discrepancy of \hat{a} and \hat{b} with respect to \hat{a} and \hat{b} , and (4) use this discrepancy to estimate the likely location of the true parameters a and b (e.g. as explained in Jonathan and Ewans [2007]). Figure 13 illustrates bootstrap estimates of parameters a and b obtained in this way together with their true locations, for a typical sample realisation from each of distributions D1-D4, using 100 bootstrap samples. Figure 13 gives an indication of the uncertainty in the pair of parameter estimates (\hat{a}, \hat{b}) for those typical sample realisations. If the modelling procedure is performing well, we would expect that the convex hull for a particular sample would include the true values of parameters for that distribution. We observe from the figure that bootstrap estimates are adequate in general, and that difference between the four data sets is less ambiguous in terms of a than b .

In summary, we find (Tables 1 and 2) that the conditional model of Heffernan and Tawn provides a relatively good characterisation of the distribution of extreme quantile Y_{10} for samples of size 1000 from four bivariate distributions for extremes, whereas the Haver model does not. This finding holds regardless of whether marginal estimation is included (Table 2) or excluded (Table 1) in the analysis. Further, conditional model parameter estimates are relatively insensitive to inclusion of marginal estimation in the modelling procedure (Figures 11 and 12). For a given sample realisation, we demonstrate good performance of a bootstrapping procedure (Figure 13) to quantify uncertainty in conditional model parameter estimates. The conditional model and Haver model are very similar in that they rely on empirical (marginal and conditional) modelling of sample data for their success. However, the conditional model structure is supported by asymptotic theory; the Haver model structure has no such motivation.

5 Applications

The four samples selected as illustrative applications for the conditional model are introduced and illustrated in Section 2. Marginal GP fits can be achieved with relatively little difficulty for storm peak H_S and T_P in each sample, as illustrated in Figures 5-8 for measured NNS data.

Following the procedure outlined in Section 3 above, the conditional model (Equation 2) is estimated for




Figure 14: Plot of residuals $\{\hat{z}_i\}_{i=1}^n$ against values $\{x_i\}_{i=1}^n$ of conditioning variate, for the measured NNS sample

peaks over threshold on data transformed to standard Gumbel marginals.

The adequacy of estimation is illustrated in Figures 14 and 15 for the measured NNS application. Figure 14 shows a scatter plot of residuals $\{\hat{z}_i\}_{i=1}^n$ against values $\{x_i\}_{i=1}^n$ of the conditioning variate. These are not obviously dependent, but do indicate skewness of the distribution of Z . Figure 15 shows the effect of varying modelling threshold on estimates \hat{a} and \hat{b} for measured NNS data. Again, optimal threshold selection is rather difficult. Figures 14 and 15 are typical of those obtained for other applications. Based on inspection of diagnostic plots such as these, it was decided to employ a 60%ile threshold for marginal modelling, and a 70%ile threshold for conditional modelling in all applications.

We are interested specifically in estimating conditional model parameters a and b , the value of the most probable 100-year storm peak H_S , H_{S100MP} and its corresponding peak period, $T_{P|H_{S100MP}}$. We assess uncertainty of these estimates using the bootstrap scheme introduced in section 4 above. Bootstrap convex hulls based on 100 resamples for each application are illustrated in Figure 16. We observe that bootstrap convex hulls contain the point estimates (\hat{a}, \hat{b}) for each application. The large extent of the convex hull for the NNS hindcast application reflects the small sample size in that case. The figure suggests that the NNS and GoM measured samples are different to one another with respect to joint extremal characteristics.

Using the conditional model we simulate realisations of exceedences of high thresholds of storm peak H_S . A typical simulation of 100 exceedences of H_{S100MP} for the measured NNS application is given in Figure 17. Using simulation we can estimate the distribution of $T_{P|H_{S100MP}}$ for a given sample. By incorporating simulation within a bootstrap scheme, we can also estimate the uncertainties in H_{S100MP} and the distribution of $T_{P|H_{S100MP}}$ for a given application. Kernel density estimates for the bootstrap distribution of H_{S100MP} for the four applications are given in Figure 18, together with point estimates in each case. Figure 19 gives corresponding estimates for the median value of $T_{P|H_{S100MP}}$.

From Figure 18 and 19 it is clear, e.g., that there is good agreement between measured and hindcast estimates for both H_{S100MP} and the median of $T_{P|H_{S100MP}}$.




Figure 15: Variation of conditional model parameter estimates \hat{a} (dots) and \hat{b} (circles) with threshold for conditional modelling, for the measured NNS sample




Figure 16: Point and bootstrap convex hull estimates for conditional model parameters a and b for the four applications. Point estimates (A-D) and bootstrap convex hulls shown corresponding to the measured NNS sample (A, circles), the NNS hindcast sample (B, dots), the measured GoM sample (C, squares), and the NNS hindcast (D, crosses).




Figure 17: Illustrative simulation from conditional model for measured NNS application, for exceedences of H_{S10} (shown as a vertical line at 15.9m).




Figure 18: Point and kernel density estimates for the bootstrap distribution of extreme quantile H_{S100MP} for each of the four applications: measured NNS (A, solid), hindcast NNS (B, dashed), measured GoM (C, dotted) and hindcast NNS (D, dashed-dotted).




Figure 19: Point and kernel density estimates for the bootstrap distribution of extreme quantile $T_{P|Hs100MP}$ for each of the four applications: measured NNS (A, solid), hindcast NNS (B, dashed), measured GoM (C, dotted) and hindcast NWS (D, dashed-dotted).

6 Conclusions and recommendations

We present an approach to joint modelling of wave climate parameters with a solid motivation from extreme value statistics. We demonstrate that the approach performs reasonably in application to simulated, measured and hindcast data. As with all extreme value analysis, estimation is relatively difficult, and threshold selection a particular issue. It is essential to demonstrate stability of the model estimates with respect to various threshold choices made. Equally, it is important to evaluate uncertainty of the modelling procedure using an approach such as bootstrapping.

The key assumption behind the conditional approach is that of a limit representation (see Heffernan and Tawn [2004] and discussion on the paper). In our simple bivariate case, this is equivalent to assuming that the residuals $Z = \frac{Y-a(x)}{b(x)}$ have the property that:

$$Pr(Z < z|X = x) \rightarrow G(z) \text{ as } x \rightarrow \infty \quad (3)$$

where G is a non-degenerate distribution, for particular choices of normalising functions $a(x)$ and $b(x)$ identified from the characteristics of the conditional distribution of $Y|X = x$. This limit form is appropriate for all existing multivariate extreme value distributions, and a large class of asymptotically independent forms. Examples do exist for which the limit representation is inappropriate, but these are somewhat obscure curiosities, at least from a present-day ocean engineering perspective. In comparison, it should also be borne in mind that the alternative assumption of multivariate regular variation is often intractable in every day practical application.

When simulating samples of sizes many times larger than the original sample used to estimate the conditional model, the set of estimated residuals $\{\hat{z}_i\}_{i=1}^n$ may be inadequate to characterise the distribution of Z , since this set is limited to the size of the original sample. In such cases, modelling the set $\{\hat{z}_i\}_{i=1}^n$ parametrically, and sampling from the fitted distribution may be more appropriate.

The conditional model and Haver model considered here are very similar in that they rely on empirical (marginal and conditional) modelling of sample data for their success. The major advantage of the conditional

model is that its structure is motivated by asymptotic theory. It would be interesting to consider whether the Haver model (for particular choices of functional forms for log-normal parameters) could be represented as a special case of the conditional approach. For example, the first and second moments of the conditional model (Equation 2) are given by $E(Y|X = x) = ax + x^b\mu_Z$ and $var(Y|X = x) = x^{2b}\sigma_Z^2$, where μ_Z and σ_Z^2 are the mean and variance of Z respectively. Standard results for the log-normal distribution yield expressions $E(Y|X = x) = \mu(x)\omega^{1/2}(x)$ and $var(Y|X = x) = \mu^2(x)\omega(x)(\omega(x) - 1)$, where $\omega(x) = \exp \sigma^2(x)$. Judicious choices of parametric forms $\mu(x)$ and $\sigma(x)$ might yield reasonable agreement between corresponding expressions for these moments (at least), possibly motivating a modified form of Haver model with improved asymptotic behaviour.

In this paper we have examined bivariate extremes only. However, the conditional model is routinely applicable in higher dimensions. Multivariate extreme value analysis concerns the characterisation, estimation and extrapolation of the joint tails of multi-dimensional distributions, and has a general relevance to meteorological-oceanographic study, far beyond the joint estimation of wave climate parameters. For example, we might wish to estimate wave climate at multiple locations in a neighbourhood, or extreme currents over a range of water depth; these require that spatial dependence of extreme values be considered carefully. The conditional model would appear to be well-suited for general application. It has become popular in recent times to use dimensionality reduction techniques such as empirical orthogonal function analysis (or principal components analysis) first to reduce the dimensionality of multivariate samples, applying extreme value analysis to data projected (for example) onto leading principal components (for example, Forristall and Cooper 1997). This approach does not model the joint extremal characteristics of the sample. It would seem that the conditional model provides a better-founded alternative.

The conditional approach provides a basis for direct simulation, not only of joint extremes (of H_S and T_P in this case) but also of functions of the variables (such as wave steepness). From a structural design perspective, it is possible to compare joint met-ocean conditions generated using a structural-response-based approach with those estimated from explicit modelling of joint met-ocean conditions.

Consideration of covariate effects is also important in extreme value analysis of meteorological-oceanographic variables (e.g. Jonathan et al. [2008]; accommodation of covariate effects by variable transformation (e.g. Eastoe and Tawn [2009]) is possible. Distributions of wave climate variables can often be partitioned according to the physical processes generating them, for example, into wind-sea and swell. The conditional model can be applied to model joint extremes of wind-sea and swell wave climate parameters. These could then be extended to define multimodal design spectra, including directionality.

Three threshold choices are necessary to estimate the conditional model in Section 5, namely a threshold to model X marginally, Y marginally and to model $Y|X$. Initial investigations suggest that marginal and conditional models are relatively stable with respect to small changes in threshold choices made. However, it would be interesting to examine more thoroughly the manner in which marginal and conditional estimation interact as a function of threshold choice, and how (in the context of 5) estimates for both H_{S100MP} and median $T_{P|H_{S100MP}}$ vary with choice of the triplet of thresholds. A further extension of the current work would be to calculate coverage probabilities for bootstrap estimates for conditional and marginal model parameters and extreme quantiles for given threshold selections.

7 Acknowledgement

The authors gratefully acknowledge discussions with Clive Anderson and Jonathan Tawn, and support of colleagues at Shell International Exploration and Production and Shell Research Ltd.

References

- E.F. Eastoe and J.A. Tawn. Modelling non-stationary extremes with application to surface level ozone. *Appl. Statist.*, 58:22–45, 2009.
- J. A. Ferreira and C. Guedes-Soares. Modelling bivariate distributions of significant wave height and mean wave period. *Applied Ocean Research*, 21:31–45, 2002.

- G. Z. Forristall and C. K. Cooper. Design current profiles using empirical orthogonal function (EOF) and inverse FORM methods. *Offshore Technology Conference OTC8267*, 1997.
- S. Haver. Wave climate off northern Norway. *Applied Ocean Research*, 7:85–92, 1985.
- J. E. Heffernan and J. A. Tawn. A conditional approach for multivariate extreme values. *J. R. Statist. Soc. B*, 66:497, 2004.
- P. Jonathan and K. C. Ewans. Uncertainties in extreme wave height estimates for hurricane dominated regions. *J. Offshore Mechanics Arctic Engineering*, 129:300–305, 2007.
- P. Jonathan, K. C. Ewans, and G. Z. Forristall. Statistical estimation of extreme ocean environments: The requirement for modelling directionality and other covariate effects. *Ocean Engineering*, 35:1211–1225, 2008.
- A. W. Ledford and J. A. Tawn. Diagnostics for dependence within time series extremes. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 65:521–543, 2003.
- A. W. Ledford and J. A. Tawn. Statistics for near independence in multivariate extreme values. *Biometrika*, 83:169–187, 1996.
- R. D. Reiss and M. Thomas. *Statistical analysis of extreme values*. Birkhauser Verlag, Basel, Switzerland, 2001.