



Technical Note: Tail behaviour of the statistical distribution of extreme storm surges

Tom Howard

Met Office, FitzRoy Road, Exeter, EX1 3PB, UK

Correspondence: Tom Howard (tom.howard@metoffice.gov.uk)

Abstract.

The tail behaviour of the statistical distribution of extreme storm surges is conveniently described by a return level plot, consisting of water level (Y-axis) against average recurrence interval on a logarithmic scale (X-axis). An average recurrence interval is often referred to as a “return period”.

- 5 Hunter’s allowance for sea-level rise gives a suggested amount by which to raise coastal defences in order to maintain the current level of flood risk, given an uncertain projection of future mean sea level rise. The allowance is most readily evaluated by assuming that sea-level annual maxima follow a Gumbel distribution, and the evaluation is awkward if we use a generalised extreme value (GEV) fit. When we use a Gumbel fit, we are effectively assuming that the return level plot is a straight line. In other words, the shape parameter, which describes the curvature of the return level plot, is zero.
- 10 On the other hand, coastal asset managers may need an estimate of the return period of unprecedented events even under current mean sea levels. For this purpose, curvature of the return level plot is usually accommodated by allowing a non-zero shape parameter whilst extrapolating the return level plot beyond the observations, using some kind of fit to observed extreme values (for example, a GEV fit to annual maxima).

This might seem like a conflict: which approach is “correct”?

- 15 Here I present evidence that the shape parameter varies around the coast of the UK, and is consequently not zero.

Despite this, I argue that there is no conflict: a suitably-constrained non-zero-shape fit is appropriate for extrapolation and a Gumbel fit is appropriate for evaluation of Hunter’s allowance.

Copyright statement. The works published in this journal are distributed under the Creative Commons Attribution 4.0 License. This licence does not affect the Crown copyright work, which is re-usable under the Open Government Licence (OGL). The Creative Commons

- 20 Attribution 4.0 License and the OGL are interoperable and do not conflict with, reduce or limit each other.

©Crown copyright 2022.



1 Introduction

Estimation of the average recurrence interval of unprecedented coastal sea level events (i.e. events with magnitude even larger than those in the tide gauge record) usually involves a characterisation of the return level plot, which is a plot of water level
25 (Y-axis) against average recurrence interval on a logarithmic scale (X-axis).

This is typically done by fitting a theoretical distribution to observed extreme values. This could mean fitting a distribution to the annual maxima, or the N-highest-ranked events per year, or all peaks over a chosen threshold. The general form of the resulting return level curve is described by the generalised extreme value (GEV) distribution (e.g. Coles, 2001), which has three parameters: location, scale and shape (μ, λ, ξ respectively).

30 Loosely speaking, the location, scale and shape can be thought of as the intercept, gradient, and curvature of the return level plot (see Howard and Williams, 2021). This technical note is concerned primarily with the *shape parameter*, which, to reiterate, can be thought of as a measure of the *curvature* of the return level plot.

Estimating the average recurrence interval of unprecedented events involves, in effect, extrapolating the return level plot beyond the domain of the observations. For this purpose, there are advantages in terms of simplicity and tractability if we fix the shape
35 parameter at zero, giving the Gumbel distribution and a straight-line return level plot. Coles (2001) and Dixon et al. (1998) counsel against this because the estimated extrapolations, and corresponding uncertainties, are sensitive to the shape parameter (as illustrated by Howard and Williams, 2021, henceforth HW21, their appendix C), and fixing it somewhat arbitrarily at zero can result in substantially underestimated uncertainties. Fixing the shape parameter affects not only the uncertainties, but also the central estimate of the extrapolation (Wahl et al., 2017, see also section 3.1 below).

40 On the other hand, Hunter (2012) proposed a simple allowance for uncertain projected future mean-sea-level rise. The method involves treating the expected number of exceedances of a given (high, rare) water level as a cost function. Weighting the sea-level rise projections by this cost function suggests an amount by which to raise coastal defences in order to maintain the current expected costs at some time in the future. Evaluation of Hunter's allowance (Hunter, 2012) in its usual form is simplified by the use of a Gumbel fit, and Woodworth et al. (2021) present evidence in favour of that fit. Furthermore, Van den
45 Brink and Können (2011) also present arguments and evidence in favour of the Gumbel fit, for extreme sea level data from tide gauges on the Dutch coast.

So, on one hand we have authors arguing in favour of accommodating curvature in the return level plot in order to extrapolate to unprecedented events, and on the other hand we have authors arguing in favour of assuming a straight line in order to simplify the calculation of a sea-level-rise allowance. In this note I suggest that these two approaches are not incompatible.

50 Martins and Stedinger (2000) note that small-sample maximum-likelihood estimators of the GEV parameters are unstable and recommend use of a Bayesian prior distribution to constrain the shape parameter. From that perspective, the use of the Gumbel distribution could be seen as an overly tight constraint on the shape parameter.



This short article presents evidence of variations in the GEV shape parameter of skew surge around the mainland coast of the UK.

55 I note that, in view of this variation, it is generally inappropriate to assume a shape parameter of zero for extrapolation of the return level curve (although, for short record lengths, the shape parameter should be appropriately constrained). However, I show that a Gumbel fit is suitable for the evaluation of Hunter’s allowance.

2 Introduction to the experiments

This article describes some experiments pertaining to geographical variations in ξ which we found whilst preparing HW21.

60 Two of the experiments were documented in HW21 and are only summarized here. The others are new and are described in detail.

2.1 Agreement in ξ from different sources

I believe that the most persuasive evidence we found of geographical variations in ξ is the discovery that the pattern of ξ diagnosed from a simulation correlates strongly with the pattern diagnosed from tide gauge data. See HW21 for details. A

65 scatter plot is shown in Fig. 1.

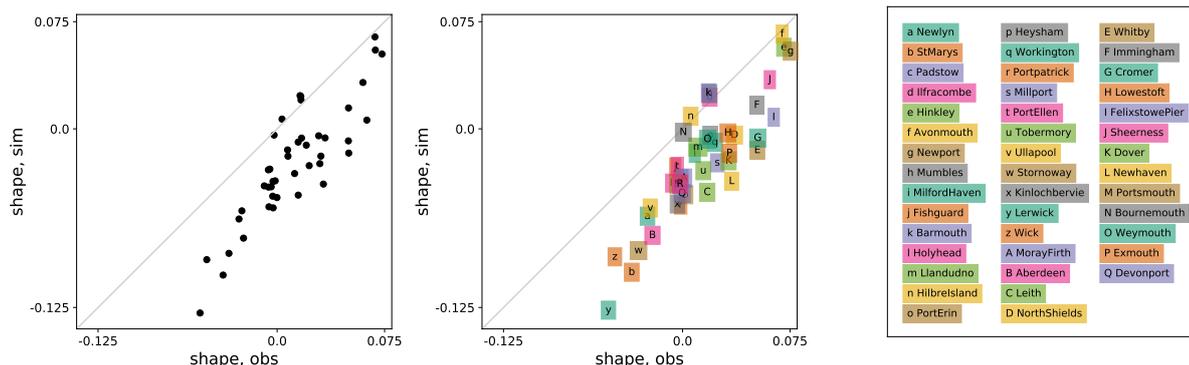


Figure 1. Left panel: Illustrating correlation between shape parameter as diagnosed from tide-gauge data (X-axis) and from a simulation (Y-axis), for sites on the UK coast. Pearson’s R coefficient is 0.86. (The right and centre panels are shown to enable identification of some of the sites.)

The simulation does not assimilate any observed data, meaning that the two patterns come from two independent data sources, and yet they correlate well. This suggests that the correlation arises due to factors which are common to both sources: essentially



the physics of the atmosphere and ocean (for example the surface momentum transfer and the bathymetry of the shelf sea) which are present in reality and represented in a numerical form in the simulation.

70 However, ξ is not the only parameter to exhibit this spatial correlation between observed and modelled values. Both μ and λ (location and scale parameters) exhibit spatial correlation. When we fit the GEV model to data, we typically find that uncertainties in λ and ξ show significant negative correlation: an underestimate in λ can be compensated by an overestimate in ξ . This raises the question of whether the spatial correlation between modelled/observed ξ could be an artifact caused by spatial correlation between modelled/observed λ somehow “printing through” to create an apparent correlation between
75 modelled/observed ξ . To control for this possibility I performed two experiments, which are described in subsection 2.2.

2.2 Break correlation

To control for the possibility of the λ — ξ compensation “printing through” into an apparent ξ (model, tide-gauge) spatial correlation via the λ (model, tide-gauge) spatial correlation and the fitting process, I performed two tests intended to break the ξ (model, tide-gauge) spatial correlation. In both tests I left the tide-gauge ξ vector unchanged. In both cases the results (de-
80 scribed below) indicate that the ξ (model, tide-gauge) spatial correlation is *not* caused by λ — ξ compensation.

2.2.1 First test

For the model data, at each site, instead of fitting GEV to the annual maxima, I fitted Gumbel to give (μ, λ) . Using (μ, λ) , I produced random Gumbel “annual maxima” for that site. I then proceeded as before (i.e. fitted a GEV distribution to that random data). If, having done this for all sites, significant ξ (random, tide-gauge) spatial correlation is still seen, it could be
85 that the ξ (model, tide-gauge) spatial correlation is an artifact of λ — ξ compensation. No such significant ξ (random, tide-gauge) spatial correlation was found.

2.2.2 Second test

This test is similar to the first test, but instead of generating random Gumbel data, I generated random GEV data, with spatially shuffled shape parameters. For the model data, at each site, I fitted GEV to the annual maxima to give (μ, λ, ξ) . For each site,
90 I used the (μ, λ) of that site, but ξ chosen at random from any site. Using (μ, λ, ξ) , I produced random GEV “annual maxima” for that site. I then proceeded as in the first test (i.e. fitted a GEV distribution to that random data). As in the first test, no significant ξ (random, tide-gauge) spatial correlation was found.



2.3 Woodworth et al. plot applied to model annual maxima

Woodworth et al. (2021) calculate Gumbel scale parameters λ for the world coastline, in order to evaluate Hunter's allowance (Hunter, 2012). The Gumbel fit is an obvious choice for this purpose: Hunter's allowance reduces the number of variables in sea-level rise projections, and although the method can be extended to accommodate the GEV, this is at the cost of reintroducing a variable (the allowance then depends on the return level of interest). Howard and Palmer (2020) demonstrate that Hunter's allowance calculations for the UK are not particularly sensitive to curvature of the return level curve, and thus a Gumbel fit is satisfactory for this purpose (see section 3).

Woodworth et al. (2021) use a simple test of the suitability of the Gumbel fit. They use the fact that, for the Gumbel distribution, the scale parameter λ is related to the standard deviation by $\lambda = \sigma\sqrt{6}\pi$ where σ is the standard deviation of the annual maxima. Thus, they use a plot of λ (by maximum-likelihood Gumbel fit to the annual maxima) against $\sigma\sqrt{6}\pi$ (where σ is the sample standard deviation) as a goodness-of-fit test.

I followed the approach of Woodworth et al. (2021) using simulated skew surges from the long control run described in HW21. Results are shown in panel (a) of Fig. 2. λ estimated from the standard deviation is within the uncertainty of the fitted estimate at most sites, but there is a consistent bias in panel (a) which is absent in the control panel (b), where Gumbel-distributed random data have been used in place of the simulated skew surge data. To quantify this bias, I evaluated the RMS difference across sites between " λ from fit" and $\sigma\sqrt{6}\pi$. To quantify the significance of this bias, I made multiple versions of panel (b) (i.e. multiple random samples using a true Gumbel distribution) and evaluated the RMS difference across sites for each version, to give a distribution of RMS difference. The RMS difference based on the real data is more than 6 standard deviations away from the mean of this random distribution, showing the statistical significance of the bias, and confirming that the simulated annual maxima are *not* Gumbel-distributed.

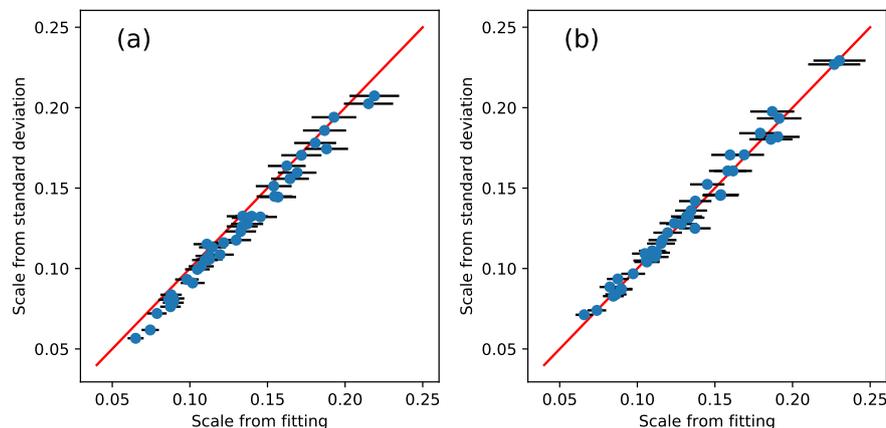


Figure 2. (a) Following Woodworth et al. (2021), we compare λ diagnosed from a Gumbel fit to 484 years of simulated annual maxima (X-axis) with $\sigma\sqrt{6}/\pi$, where σ is the standard deviation of the 484 annual maxima. The two should be the same in the case of the annual maxima conforming to a Gumbel distribution. The horizontal error bars represent 95-percent uncertainty in the fitted λ . The red line represents a ratio of 1. Each point represents the site of a tide gauge on the coast of the UK. (b) as (a), but instead of simulated annual maxima, at each site we use 484 random variates drawn from the Gumbel distribution, using the (μ, λ) determined by the fit in (a).

2.4 Tests following Van den Brink and Können (2008)

Van den Brink and Können (2008) (henceforth VdBK08) devised a sophisticated goodness-of-fit test that emphasises the quality of fit at the most extreme value (the “outlier”) from each site. They developed what might be called a “standardised” quantile-quantile (QQ) plot¹ showing the outlier at each site, where “standardised” refers to removing the effects of different GEV parameters and record lengths at different sites. The approach is described rigorously in VdBK08. A brief informal expression of an exactly equivalent procedure follows.

Consider the following concept.

120 PIT If the cumulative distribution function (CDF) of a random variate R is known, then a sample $R_i, i = 1 \dots n$ of n values from that distribution can be transformed to a sample $U_i, i = 1 \dots n$ whose expected distribution is standard-uniform. This is sometimes referred to as the probability integral transform (PIT, see for example Folland and Anderson, 2002). If the inverse CDF is known, the PIT can also be used to generate a non-uniformly distributed sample such as $R_i, i = 1 \dots n$ from a standard-uniformly distributed sample such as $U_i, i = 1 \dots n$.

125 We use this concept in the following procedure.

¹ A quantile-quantile plot (sometimes “quantile plot” (e.g. Coles, 2001)) typically compares the quantiles (water levels) from two different sources, for example observations on one axis and a statistical model on the other axis.



Step 1 Suppose that the distribution of a sample $Y_i, i = 1 \dots n$ of n annual maxima from a given site is generalised extreme value (GEV).

Step 2 If the CDF of Y is known, we can transform this sample using the PIT to a sample $U_i, i = 1 \dots n$ whose expected distribution is standard-uniform. Knowing the CDF of the GEV means knowing its parameters. So, this transformation depends on the parameters.

Step 3 The “outlier” of this transformed sample, $M = \max(U_i, i = 1 \dots n)$ has its own distribution: $\Pr(M < x) = x^n$ (notice that this depends on n).

Step 4 Since its CDF is known, a sample $M_j, j = 1 \dots m$ of outliers from m different sites can also be transformed to a standard uniform distribution, $U_j, j = 1 \dots m$, using the PIT. (n need not be the same at every site).

(Reminder: n is the number of annual maxima at a given site. m is the number of sites considered.) We can plot the sorted sample $U_j, j = 1 \dots m$ against its expected (i.e. uniformly-distributed) values. Departures from this expectation will arise due to sampling uncertainty, resulting in departures from the idealised straight line of gradient one and intercept zero.

VdBK08 apply this procedure to a situation where the CDF of Y at each site is *not* known, but rather estimated by fitting a GEV (or Gumbel) to the sample $Y_i, i = 1 \dots n$ at that site. Now departures from the idealised line will arise not only from sampling uncertainty, but also from poor fitting. Thus the departure from the idealised line becomes a measure of goodness-of-fit, with the emphasis on quality of fit at the outlier at each site. This plot can be thought of as a “standardised” probability-probability (PP) plot², where “standardised” refers to removing the effects of different GEV parameters and record lengths at different sites. VdBK08 perform one further transformation. Instead of the plot described above, they take $-\log(-\log(\cdot))$ of both axes. This emphasizes the departures at the high-end extremes of the plot, making their plot more like a “standardised” QQ plot, where again “standardised” refers to removing the effects of different GEV parameters and record lengths at different sites. It is also known as a Gumbel plot, $-\log(-\log(\cdot))$ being the inverse CDF of the Gumbel distribution.

VdBK08 show that when unconstrained GEV is fitted to synthetic Gumbel-distributed maxima, the additional free parameter results in over-fitting; the variability in $U_j, j = 1 \dots m$ is seen to be too small.

Among other applications, Van den Brink and Können (2011) used the procedure specifically to help choose between GEV-fitting and Gumbel-fitting the distributions of simulated and observed annual maximum sea level at 19 sites on the Dutch coast. They fixed the shape parameter and found an optimum value of -0.005. They prefer Gumbel-fitting to unconstrained GEV-fitting for this application, arguing that the uncertainties in extrapolation are smaller.

I followed the procedure described to make both types of plot (PP and QQ) for two different sets of annual maximum skew surge at the locations of UK mainland tide gauges. The two different sets are tide gauge data and data from a numerical simulation (for details see Howard and Williams, 2021). The resulting plots are shown in Fig. 3.

² A probability-probability plot (sometimes “probability plot” (e.g. Coles, 2001)) typically compares the probability of an event as estimated from two different sources, for example estimated empirically from observations on one axis and estimated from a statistical model on the other axis.

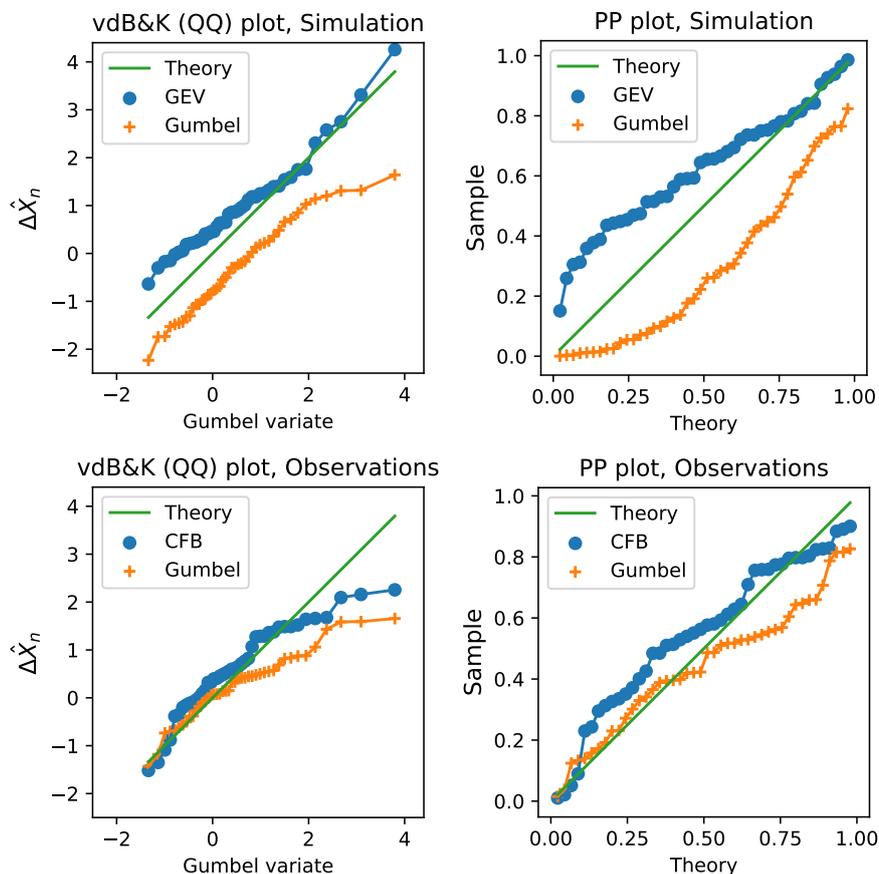


Figure 3. Left-hand panels show plots following VdBK08 and right-hand panels show PP plots as described above, for skew surge at the sites of UK mainland tide gauges (one plotted point per gauge). Orange cross symbols relate to a Gumbel fit to annual maxima. Blue circles relate to unconstrained GEV fits to annual maxima (simulation, top panels) and constrained GPD fits to peaks over a threshold following the Coastal Flood Boundary Conditions (CFB) guidance (Environment Agency, 2018) method (observations, bottom panels). Each point represents an outlier (the largest annual maximum in the record) at a gauge.

As in VdBK08 (their figure 3), each plot shows a comparison of two different types of fitting. The top panels compare unconstrained GEV fits (blue circles) to simulated annual maxima with Gumbel fits (orange crosses) to the same annual maxima. These results suggest that allowing for some variation in the shape parameter is preferable to Gumbel fitting, presumably because Gumbel fitting represents too tight a constraint on the shape, for these data.

160 In the bottom panels, the sophisticated constrained generalised Pareto distribution (GPD) fit to peaks over a threshold (POT) as used in the Coastal Flood Boundary Conditions (CFB) guidance (Environment Agency, 2018) was applied to observed skew surge extremes, and the blue circles relate to that fit. The orange crosses again relate to simple Gumbel fits to annual maxima.



This is a less fair comparison than that shown in the top panels, because the constrained GPD fit to POT takes advantage of more data. Nevertheless, the constrained GPD fit allows for a non-zero shape parameter whereas the Gumbel fit does not, and again we do not see any clear support for preferring the Gumbel fit.

One possible criticism of this test is that, since a surge event will typically effect several sites, the data are not independent. A crude fix is to miss out closely-neighbouring tide gauges, and to use results from every second tide gauge, every third tide gauge, etc. The results of doing so are shown in the Howard and Williams (2021) review material available online (see bibliography for a URL). Again, these results do not provide any clear support for preferring the Gumbel fit.

170 3 Discussion

All of the above experiments illustrate that the shape parameter of sea level annual maxima around the UK is not, in general, zero. Thus it is appropriate to use a (suitably constrained) fit accommodating non-zero shape for the purpose of quantifying the return level curve and extrapolating to quantify levels (and uncertainties) of unprecedented events at these sites.

However, Howard and Palmer (2020) (working with still water level rather than skew surge) have shown that curvature in the return level plot gives variations of less than 5 cm (absolute, and less than 6% relative) in Hunter's allowance for UK coastal sites based on RCP8.5 for 2100, and this contribution to uncertainty in the allowance is less important than the contribution from the uncertainty in the present-day return level curves (see Fig. 4 here, and Howard and Palmer, 2020, their section 5.3). Thus, it is reasonable to use the Gumbel distribution for a simple evaluation of Hunter's allowance, at least for sites on the UK coast. Note, however, that curvature in the return level plot implies that an appropriate allowance may be dependent on the return period of interest (Buchanan et al., 2017).

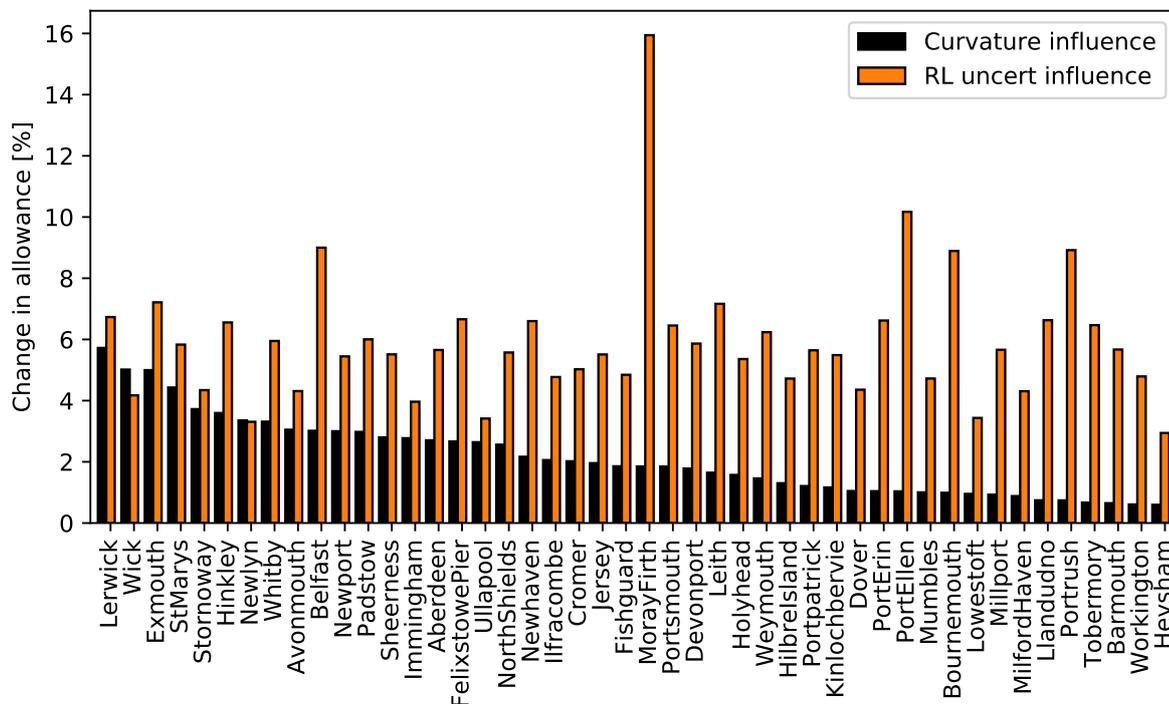


Figure 4. Following Howard and Palmer (2020), we show the sensitivity (black filled bars) of Hunter’s allowance to curvature of the return level plot of still water level (in other words, sensitivity to the departure of the still water extremes from a Gumbel distribution). Also shown for comparison is the sensitivity to uncertainty in the return level curve. For full details of the method see Howard and Palmer (2020). Sites are ordered by curvature sensitivity (not geographical location).

3.1 Extrapolation

Here are some simplified examples illustrating the effect of constraint when extrapolating the return level curve beyond the observational record. I compare three different parametric fits to observed annual maximum skew surge at UK coastal sites:

1. GEV fit unconstrained
- 185 2. GEV fit constrained to give a range of shape parameters comparable to CFB2018.
3. Gumbel fit

Maximum likelihood estimation (MLE) is used in all cases (Generalised MLE (Martins and Stedinger, 2000) for case 2). CFB2018 used a normal prior on the shape with mean 0.0119 and standard deviation 0.0343. Applying this to only the annual maxima gives a much smaller range of shapes than that found by CFB2018. This is because the annual maxima represent a
 190 smaller data set than the peaks over threshold used by CFB2018, and this smaller quantity of data allows the prior to dominate.



To compensate for this, I increased the width of the prior until I produced a similar range of shapes to CFB2018. My prior has mean 0.0119 and standard deviation 0.062. Using only the annual maxima is not to be recommended in general, because better results can be obtained using more data (for example N-largest or peaks-over-threshold). Nevertheless, I take the constrained GEV fit to be “truth” for the purpose of this illustration, and identify anomalies relative to that. Anomalies are shown in fig. 5.

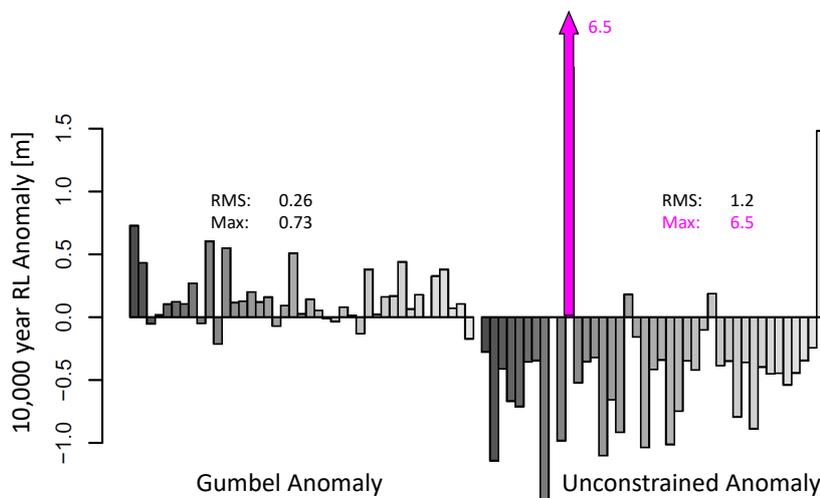


Figure 5. Anomalies in the ten thousand year return level of skew surge as estimated from annual maxima at 44 UK coastal sites by maximum likelihood estimation using a Gumbel fit (left) and an *unconstrained* GEV fit (right). Anomalies are relative to a *constrained* GEV estimate (see main text). The bar shown in pink (Hinkley Point) exceeds the limits of the plot.

195 The most serious anomaly in the unconstrained GEV fit is at Hinkley Point in the Bristol Channel, where only 26 skew surge annual maxima from the tide-gauge record are included in the fit. Figure 6 illustrates the three different fitted return level curves at that site. Also shown is an example site (Bournemouth, 18 annual maxima), where the unconstrained shape parameter is negative.

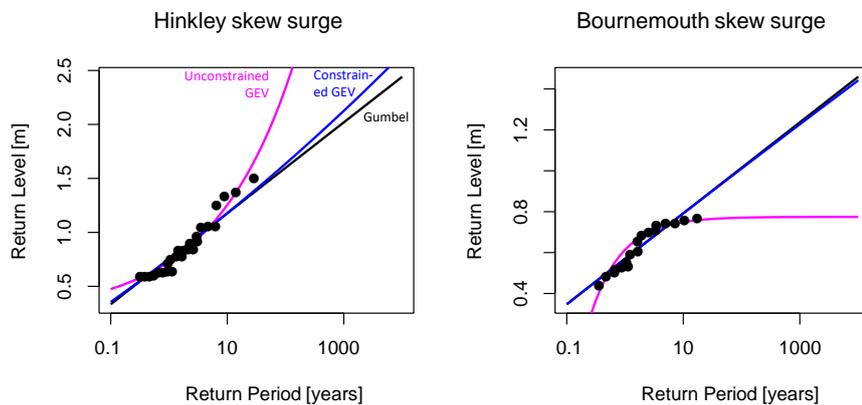


Figure 6. Illustrating a well-known issue with unconstrained GEV fit by MLE to annual maxima from a short record. The three different fits are described in the main text.

We have seen in previous sections that the shape parameter of UK skew surges varies spatially and, consistent with this, we have some confirmation (in section 2.4) that an appropriately constrained GEV fit is preferable to a Gumbel fit for UK skew surges. Figures 3, 5 and 6 all show that the Gumbel fit is at odds with an appropriately constrained GEV fit, but on the other hand, Figs. 5 and 6 show that the Gumbel fit is much safer than an unconstrained GEV fit to short records.

4 Conclusions

In summary, at least for sites on the UK coast, a non-zero shape parameter should be accommodated at the fitting stage for the purpose of extrapolating the storm surge return level curve. However, fitting with an unconstrained shape parameter to short records is not advisable, as it is liable to give larger errors than the over-constraint inherent in a Gumbel fit.

Also, it is reasonable to use a Gumbel fit for evaluation of Hunter's allowance.

Data availability. Contains Environment Agency information ©Environment Agency and database right

Author contributions. TH performed the tests and wrote the text. Please also see the acknowledgements.

Competing interests. The author declares no competing interests.

<https://doi.org/10.5194/os-2022-14>
Preprint. Discussion started: 18 March 2022
© Author(s) 2022. CC BY 4.0 License.



Acknowledgements. This work was supported by the Met Office Hadley Centre Climate Programme funded by BEIS and Defra. Thanks to Simon Williams and Jenny Sansom for help with the CFB data.



References

- Buchanan, M. K., Oppenheimer, M., and Kopp, R. E.: Amplification of flood frequencies with local sea level rise and emerging flood regimes, *Environmental Research Letters*, 12, 064 009, 2017.
- 215 Coles, S.: An introduction to statistical modeling of extreme values, Springer, 2001.
- Dixon, M. J., Tawn, J. A., and Vassie, J. M.: Spatial modelling of extreme sea-levels, *Environmetrics: The official journal of the International Environmetrics Society*, 9, 283–301, 1998.
- Environment Agency: Coastal Flood Boundary Conditions for the UK: update 2018, Tech. rep., <https://environment.data.gov.uk/dataset/84a5c7c0-d465-11e4-b0bd-f0def148f590>, Available to download from <https://environment.data.gov.uk>, 2018.
- 220 Folland, C. and Anderson, C.: Estimating changing extremes using empirical ranking methods, *Journal of climate*, 15, 2954–2960, 2002.
- Howard, T. and Palmer, M. D.: Sea-level rise allowances for the UK, *Environmental Research Communications*, 2, 035 003, <https://doi.org/10.1088/2515-7620/ab7cb4>, 2020.
- Howard, T. and Williams, S. D. P.: Towards using state-of-the-art climate models to help constrain estimates of unprecedented UK storm surges, *Natural Hazards and Earth System Sciences*, 21, 3693–3712, <https://doi.org/10.5194/nhess-21-3693-2021>, (online review material available at <https://nhess.copernicus.org/preprints/nhess-2021-184/nhess-2021-184-AC1-supplement.pdf>), 2021.
- 225 Hunter, J.: A simple technique for estimating an allowance for uncertain sea-level rise, *Climatic Change*, 113, 239–252, <https://doi.org/10.1007/s10584-011-0332-1>, 2012.
- Martins, E. S. and Stedinger, J. R.: Generalized maximum-likelihood generalized extreme-value quantile estimators for hydrologic data, *Water Resources Research*, 36, 737–744, 2000.
- 230 Van den Brink, H. and Können, G.: The statistical distribution of meteorological outliers, *Geophysical research letters*, 35, 2008.
- Van den Brink, H. and Können, G.: Estimating 10000-year return values from short time series, *International Journal of Climatology*, 31, 115–126, 2011.
- Wahl, T., Haigh, I. D., Nicholls, R. J., Arns, A., Dangendorf, S., Hinkel, J., and Slangen, A. B.: Understanding extreme sea levels for broad-scale coastal impact and adaptation analysis, *Nature communications*, 8, 1–12, 2017.
- 235 Woodworth, P. L., Hunter, J. R., Marcos, M., and Hughes, C. W.: Towards reliable global allowances for sea level rise, *Global and Planetary Change*, 203, 103 522, <https://doi.org/https://doi.org/10.1016/j.gloplacha.2021.103522>, 2021.