

Tom Howard 9 Sep. 21
Author response to:

Discussion of Towards using state-of-the-art climate models to help constrain estimates of unprecedented UK storm surges

by Tom Howard and Simon David Paul Williams

Discussants: Eleanor D’Arcy and Jonathan Tawn

August 27, 2021

Thank you both very much for your thorough and interesting review. I have pasted in the whole discussion in dark orange font colour; my responses are shown in blue.

1 Overview

This paper proposes using climate model simulations to aid with still water level return level estimation and address problems that arise with the statistical approach when tide gauges have short record lengths. The model they present is the HadGEM3-GC3-MM to generate a dataset of 483-year present-day storm surges at sites on the UK tide gauge network. They compare the skew surge simulations to using only observations when fitting extreme value models by estimating parameters. The spatial distribution of the parameters for each dataset are generally well correlated. However, there is a negative bias in the simulation approach in the shape parameter estimate of the generalised Pareto distribution (GPD); the authors discuss this in detail and suspect it is due to a pitfall in the shelf sea model (CS3).

The paper also investigates the interaction between skew surge and peak tide at Sheerness. They study the effect that changes in timing between atmospheric forcings and tide has on skew surge. Additionally, they review the independence assumption of skew surge and peak tide used in the JPM. Using their model simulations, they show that extreme skew surges are more likely to occur on neap tide - this agrees with the results of Williams et al. (2016) (supplementary material).

2 General Comments

The paper is well written, in the proceeding sections we have discussed parts of the paper that pose interesting areas for future research and noted some technical corrections. The results are well presented and the figures well explained, giving a clear justification for the proposed model. The appendix provides strong support for ideas mentioned in the main paper. The authors have recognised potential downfalls with different aspects of the model and presented some initial investigation into these (for example, the discussion of why the shape parameter is more negative for simulations on pg. 15/16).

3 Specific Comments

3.1 Comparison of Model and Observed Data

Figure B1 shows some major departures between the model and observed data across the distribution of skew surges but particularly in the tails. In no sites is the model giving as high quantiles as the observed data. However, these departures have a systematic feature which is consistent over spatial regions, e.g., south-west and north-west UK. This suggests that it should be possible to account for these departures through a smooth spatial function which maps the differences in quantiles between the observations and model data. With this adjustment it is possible that the currently identified under-estimation may be corrected before making the tail based GEV/GPD comparisons you draw.

Have added the following to the description accompanying Figure B1:

“Figure B1 shows some major departures between the model and observed data across the distribution of skew surges, but particularly in the tails. The model does not give higher quantiles than the observed data at any site.”

And added the following to the main text:

“The model does not give higher quantiles than the observed data at any site.”

Have added a section “Suggestions for further work”, paraphrasing your comment in the following text:

“The model/observation departures seen in Fig. B1 have a systematic feature which is consistent over spatial regions, e.g., south-west and north-west UK. This suggests that it should be possible to account for these departures through a smooth spatial function which maps the differences in quantiles between the observations and model data. With this adjustment it is possible that the currently identified under-estimation may be corrected before making the tail-based GEV/GPD comparisons shown here.”

On first reading your suggestion, I thought: “that would only correct the location and scale parameters, which we are suggesting be taken from the observations anyway --- it wouldn’t make sense to correct the model shape parameters using the shorter observational records and then argue that the observational shape parameters should be replaced with those of the model”. But then I remembered about the scale-shape compensation which can occur at the fitting stage. Is that the reason for your suggestion?

3.2 Shape Parameters with Similar Spatial Patterns

One of your exciting findings is that the spatial pattern of the shape parameter estimates is similar for the CFB2018 estimates from observed data and your estimates from the model, but with a systematic bias between them. This suggests using an alternative to the CFB2018 approach by explicitly exploiting this finding. Let $\xi_{obs}(x)$ and $\xi_{model}(x)$ be the shape parameters for the observed and model data respectively, for all gauged sites x . What you are saying is that you believe in the spatial variation of $\xi_{model}(x)$ but not its mean value. So you believe that $\xi_{obs}(x) = \xi_{model}(x) + \xi_{bias}$, where ξ_{bias} is a fixed constant that does not vary over x . Your estimates indicate that $\xi_{bias} > 0$. Fixing estimates of $\xi_{model}(x)$ for all x and fitting the function of $\xi_{obs}(x)$ over x now means only one parameter, ξ_{bias} , needs estimating. This could lead to substantial reductions in the uncertainty of $\xi_{obs}(x)$ estimates over x .

I like that idea a lot! Have incorporated it in the main text. Hope that is OK. Have acknowledged your help in the acknowledgements section.

Line 344: have removed the phrase “without the subjectivity of a prior” and added the following:

“Given the need for some kind of constraint on the shape parameter when fitting observational records, use of shape parameters from a long simulation holds the promise of reducing uncertainties. For example, if we assume that the

model-diagnosed spatial pattern of shape parameters is correct but uniformly biased by a scalar ξ_{bias} (which does not vary over sites), $\xi_{\text{true}}(x) = \xi_{\text{model}}(x) + \xi_{\text{bias}}$, where x is a vector of sites, then we can use the observations from *all sites* to estimate the one scalar parameter ξ_{bias} . This could lead to substantial reductions in the uncertainty of ξ_{true} estimates over x .”

3.3 Penalised Likelihood

The shape parameter estimates in Figure 2 (d) based on the 483 years of model data show some site-to-site variations which are more pronounced than the broader smooth variations across coastlines. This suggests that they would also benefit from the penalty-based approach used in CFB2018 work.

Thanks, have paraphrased this in the new “Suggestions for further work” section.

At a few points you state that the prior/penalty is subjective and that this is a disadvantage. Yet you also point that the smooth pattern of the shape parameter estimates this gives for the observed data agrees well with the similar unpenalised estimates using model data. You say this is a really positive feature of the model data, we would also take this as supporting the value of the process of creating penalised estimates from the observed data. The penalty is giving something meaningful, so the effect of the claimed “subjectivity” is positive for the observed data analysis.

I completely agree. I’m sorry if this did not come across clearly in the draft. Indeed, in some other experiments (not discussed in the paper) I tried unconstrained GEV fits to the annual maxima from the observations. This spoiled the strong correlation with the model-diagnosed shape parameters, so yes, completely agree that the CFB2018 penalisation process adds value over an unconstrained fit.

In the draft, we have the following text:

This strong correlation between the two spatial patterns of shape parameter diagnosed from independent sources (i.e. our model simulation and the tide-gauge data) is remarkable. It both supports the spatial pattern of the shape parameter as a real, physically-determined phenomenon (as opposed to a statistical artefact), **and gives further credibility to both the CFB2018 approach** and our model. [bold font is not in the draft].

The prior that was selected in the CFB2018 work was not subjective in the traditional sense of a subjective prior in Bayesian methods. It was actually a data-based prior which corresponds to an empirical Bayesian prior, using all the information that separately estimated shape parameters for UK skew surge provide. The effect of this was simply to move shape parameter estimates more towards the UK average, with the larger changes coming for sites with shorter record lengths.

Thank you for your guidance. I didn’t previously understand that difference in the definition of subjective vs empirical. Have added your description to the draft as follows, and removed all references to “subjectivity”

“CFB2018 employed a data-based prior using all the information that separately estimated shape parameters for UK skew surge provide. The effect of this was simply to move shape parameter estimates more towards the UK average, with the larger changes coming for sites with shorter record lengths.”

3.4 Investigation of Interaction between Tide and Skew Surge

We are really pleased to see the use of numerical models to explore systematically the widely claimed property that skew surges are independent of their associated peak tidal values. The work in Section 5 where you explore the timing effects of skew surge events relative to tides is very illuminating. It adds greatly to the existing empirical evidence for this feature at Sheerness. It would be interesting to have your opinions about what interactions are expected elsewhere in the UK given that the CFB2018 estimates over UK sites all assume independence. Since this analysis is based on simulations from HadGEM3-GC3-MM, this presents a physically-based justification for dependence between skew surge and peak tide at Sheerness. It suggests future research is required to investigate this. It may be that interaction occurs to some extent everywhere, but at a practical level it is not important apart from some locations.

Agreed. Have added a sentence to the “Suggestions for further work” section. For what it’s worth, my guess is that interaction is stronger at Sheerness than elsewhere, and is associated with the surge and tide travelling a long way together (all the way down the east coast) in shallow water. But that is just a guess.

In current work we have been investigating this feature empirically at a limited number of sites. Here we report some of the methods and findings using data from the tide gauges at Sheerness and Heysham (located close to Workington which you analyse). We study data from Heysham in 1964-2016, of which 17.5% is missing, and 1980-2016 at Sheerness, where 9.1% is missing.

We define extreme skew surges as exceedances of the 0.95 quantile at each site. Figure 1 shows scatter plots of extreme skew surges against their associated ranked peak tide; these are ranked so that 1 corresponds to the smallest observation and n is the largest, where n is the total number of observations. If the two components were independent, we expect extreme skew surges to be uniformly distributed over ranks. We test this using a Kolmogorov-Smirnov test for uniformity; the p value at Heysham is 0.0224 whilst at Sheerness it is 1.40×10^{-7} . Clearly the p value at Sheerness is much smaller, and provides statistical evidence that extreme skew surges are not independent of peak tide. However, at Heysham, there is not sufficient evidence to reject the claim of independence at the 0.01 significance level. This is clear in Figure 1, where more extreme skew surges occur on lower tides at Sheerness - this agrees with your findings and those of Williams et al. (2016).

We investigate this further by looking at skew surge and peak tide dependence on a month-by-month basis, using ideas from Williams et al. (2016). Figure 2 compares the distribution of peak tides associated




Figure 1: Extreme skew surge observations against ranked peak tide at Heysham (left) and Sheerness (right).




Figure 2: Monthly distributions of peak tides at Sheerness in February, May, August and October. The probability density function of all peak tides (black) and peak tides associated with extreme skew surge (red) are interpolated onto each distribution.

with all skew surges and the distribution of peak tides associated with extreme skew surges for February, May, August and October. If peak tide and skew surge are independent, these two distributions should be the same, up to sampling variation. We estimate the probability density function (pdf) using a Gaussian kernel density estimate, and use an Anderson-Darling test to check if peak tides come from the same distribution as peak tides associated with extreme skew surges. Figure 2 highlights how the dependence between skew surge and peak tide is changing with the time of year; the distribution of peak tides and the peak tides associated with extreme skew surge are most different in May and least different in February. In May, the mode of the distribution of peaks tides associated with extreme skew surges has shifted to a lower value than the distribution of all peak tides. Results from the Anderson Darling test for every month tell us there is insufficient evidence to reject the null hypothesis that peak tides and the peak tides associated with extreme skew surge come from the same distribution in February, March, September, November and December. In these months, we conclude it is reasonable to assume skew surge and peak tide are independent. In the remaining months, we find sufficient evidence to suggest the two components are dependent. We believe this poses a really interesting area for further research.

In D’Arcy et al. (2021) we fit a GPD (see Coles (2001) for details) to extreme skew surges that accounts for seasonal variations, since they are more extreme in the winter. Our results show this is an improvement on a standard GPD fit, where skew surges are assumed to be independent of peak tide. Here, we investigate adding a covariate of peak tide into our skew surge model to account for the dependence found at Sheerness. D’Arcy et al. (2021) defines extreme skew surges as exceedances of the monthly 0.95 quantile. Non-stationarity is accounted for in the scale parameter of the GPD through a daily covariate $d = 1, \dots, 365$:

$$\sigma_d = a + b \sin\left(\frac{2\pi}{365}(d - \phi)\right) \quad (3.1)$$

for $a, b, \phi \in \mathbb{R}$ parameters to be estimated. Note that the shape parameter of the GPD is fixed across months. To account for the dependence between skew surge and tide, we now also consider the following parameterisation on the scale parameter:

$$\sigma_{d,t} = a + b \sin\left(\frac{2\pi}{365}(d - \phi)\right) + ct \quad (3.2)$$

where $c \in \mathbb{R}$ is another parameter to be estimated, and t the associated peak tide observation. We fit a GPD with both (3.1) and (3.2) formulations to extreme skew surges at Heysham and Sheerness. The Akaike information criteria (AIC), frequently used for model selection, suggests that formulation (3.1) gives a better model fit at Heysham, i.e., an independence conclusion is supported by this analysis. Whereas, AIC suggests that the parametrisation in equation (3.2) yields a better fit at Sheerness. By ordering peak tide observations from smallest to largest and calculating $\sigma_{d,t}$ at its winter peak ($d = 365$) for each value, we observe a 15% reduction in the scale parameter as tides increase at Sheerness, which corresponds to an equal percentage reduction in the skew surge quantiles for

excesses of the December skew surge threshold. We also compare the model fits at each site using a Likelihood Ratio Test. At Heysham the p value is 0.657 which provides insufficient evidence to reject the null hypothesis, that is the simpler model (in equation (3.1)) is sufficient. Whereas, at Sheerness we get a much smaller p value of 0.059, which provides statistically significant evidence at the 0.1 significance level to reject the null hypothesis and conclude that the more complex model is required. This highlights the importance of accounting for skew surge and peak tide dependence at sites where the independence assumption is not justified.

This is very interesting. Thank you for including it in your review. In the draft I have added a citation to your forthcoming publication.

3.5 Ungauged Sites

We feel you do not do full justice to your developments given the focus is on comparisons made at gauged sites with long and trustworthy records. The real value in modelled data is the ability to give estimates at other sites. This could be the focus of a natural follow up paper.

Have added a sentence in the “Suggestions for further work”

3.6 Other Comments

1. In the abstract, you say “results suggest an event of this magnitude has an expected frequency of about 1 in 500 years at [Sheerness]” when referring to the North Sea floods of 1953. In the paper, the only result to show this is presented in Section 6.2: “the fact that the 483-year surge-only simulation produces more than one event of comparable magnitude to the simulated 1953 event suggest the return period of the 1953 atmospheric forcing is less than 483 years.” If this is the justification, it would be better to more formally quantify this using your statistical models.

I can see several different ways to approach that. In the preceding sections we have advocated using only the shape parameter from the simulations. Using the observational location and scale parameters (as in Fig. 2 panels (a) and (b)) and the simulation-diagnosed shape parameter (panel (d)), the range of the Wadey skew surges as shown in our figure 7 correspond to return periods ranging between about 650 and 2000 years. (Note to self: code_ABR). Having said that, I feel that this level of detail is not appropriate in this preliminary “Towards using state-of-the-art climate models...” type of publication, so I have cut the reference to the expected frequency out of the abstract, and cut the corresponding short paragraph from the main text.

2. In the Introduction (line 51) you list assumptions required to fit extreme value models as ‘events are effectively random and statistically independent of each other.’ But what about events being identically distributed (or stationary), it is clear that tide and storm surge (or skew surge) are not stationary as both process exhibit seasonality. Instead of “effectively random” it would be more mathematically correct to say “stochastic.” Extreme value methods also handle dependence, so “independent of each other” is not formally required.

Rephrased as follows:

The statistical models which are fitted to the observational data in order to infer the levels of unprecedented extremes are supported by mathematical arguments which may require assumptions such as the assumption that the events are stochastic. We know that the real-world events are deterministic, and furthermore may be auto-correlated over a range of timescales. Such auto-correlation can be accounted for within the statistical framework, for example by the use of an extremal index^{~\citep{Tawn1992estimating, Batstone2013UK}}. Alternatively, a physically-based numerical model has the potential to directly address both determinism and auto-correlation by simulating them.

3. On line 84 climate change is discussed. Tide gauge observations will exhibit an approximately linear mean trend due to sea level rise, it is unclear whether this has been removed before comparison with the simulations in Section 4. It is likely that removing this change will not change the results significantly, but it is important to remove this non-stationary effect before fitting extreme value models.

Have added this phrase:

“The trend due to sea level rise can be seen in the tide gauge observations and was carefully removed before making a statistical fit to the extremes. For details see CFB2018. Our numerical model of the shelf sea does not include any change in mean sea level.”

4. Figure C1 shows the reduction in uncertainty on shape parameter when record lengths are increased in the tide gauge network, but it would be nice to see this for more record lengths that are equally spaced (say 10 to 500 years in increments of 10 years) to really highlight this - with measures of uncertainty as in Figure C1 (d).

Good idea. Done. Thanks.

4 Technical Corrections

1. Table 1 gives a list of useful acronyms and symbols, it would help to include what ‘GC3’ and ‘MM’ stand for in ‘HadGEM3-GC3-MM’ as it is not mentioned in text.

Thank you for pointing this out. Have added it to Table 1.

2. Equation (2) has a surplus close bracket. Corrected, thanks.
3. Equation (6), should this derivative be evaluated at $L = \log(u)$ rather than $y = u$ since y is the return level (as in equation (1))?

No, I don’t think so. However I acknowledge that, in equation 6, the point on the RL curve at which to evaluate the gradient is specified in a non-standard way: in terms of the ordinate (y) instead of the usual specification in terms of the abscissa (L). Your suggested phrase “evaluated at $L = \log(u)$ ” does not mean the same thing. Please see the uploaded supplement **equation6_revisited.pdf** for further explanation.

Note to self: upload eq 6 supplement.

4. Figures 2 (c) and (d) would benefit from having 95% confidence intervals on shape parameters using the model data.

Have added these, and commented in the text on the added value in terms of certainty that the CFB method provides compared to a simple GEV fit to annual maxima --- this is well-illustrated by the confidence intervals shown.

5. Line 283: We think the wrong figure is referenced here.

Please note this is figure E.1 in the CFB technical publication (“**their** figure E.1”, as stated in the text). In the copy I downloaded from here:

https://assets.publishing.service.gov.uk/media/603652cce90e0740b7caac9d/Coastal_flood_boundary_conditions_for_the_UK_2018_update_-_technical_report.pdf

figure E.1 is on page 64, titled “Estimated shape parameter for the UK tide gauges.”

6. Line 319: "They usedtime-series." This has nothing to do with how CFB2018 estimate the shape parameters and should be cut. It only links to the derivation of SWL return levels.

Have cut these two sentences as advised.

7. Line 349: You hint that the reason for the disagreement between the methods could be due the short observed data series. But Figure B1 shows major departures between the observed data and the model data in the body of the distributions to a level that could in no way be due to short observed series and must be that the model data does not reproduce well the observed data. Linked to this, in discussing Figure 1 you say the agreement at both sites is "excellent". With the amount of data at Sheerness it is clear that there are major disagreements that cannot be accounted by sampling variations.

Have reworded that paragraph to be less bullish:

"Figure 1 shows good model vs observations agreement at Workington, and even these two sites alone illustrate that our modelling system is able to simulate unprecedented skew surge events (i.e. events of a magnitude not found in the tide-gauge record). However, the quality of agreement shown at Workington is not exhibited everywhere. Empirical return level plots of skew surge for a set of 44 tide gauge locations around the UK coastline are shown in the appendix in Fig. B1. This gives a qualitative, visual sense of the realism of the model in terms of the simulated extremes. The good agreement at Workington can be contrasted with the poor agreement at, for example, Newlyn or Aberdeen, where the simulated extremes are negatively biased relative to the corresponding observations. The model does not give higher quantiles than the observed data at any site. We argue (below) that the simulation may nevertheless be able to add value to estimations of unprecedented events, even where a bias exists."

8. Line 421: You suddenly mention a "kernel" to spread the duration of events but fail to provide any information. A little detail would be helpful here.

I have added the following sentences explaining the choice and purpose of the kernel:

"The kernel was designed to represent the important features of the RACMO-driven surge-only simulation, i.e., the approximate duration and shape of the time series plot. The purpose of convolution with the kernel is to identify those events which correlate well (in terms of their time series plot) with the RACMO-driven simulation, in other words, events which not only produce a large surge, but are also of comparable duration to the RACMO-driven simulation. The kernel was not used to modify events, but simply to identify significant ones."

9. Figure E2: A description of the different points would be useful.

Have added the following text: "The blue points show simulated skew surge against simulated astronomical high water for the 16 extreme atmospheric events with timing adjusted such that the event coincided with a simulated spring tide and a simulated neap tide, as described above. The artificial case of no tide is also shown. Grey points are as in Williams (2016)."

References

Coles, S. G. (2001). *An Introduction to Statistical Modeling of Extreme Values*. Springer, London.

D'Arcy, E., Tawn, J., Joly-Laugel, A., and Sifnioti, D. E. (2021). Accounting for seasonality in extreme sea level estimation (*in preparation*).

Williams, J., Horsburgh, K. J., Williams, J. A., and Proctor, R. N. (2016). Tide and skew surge independence: New insights for flood risk. *Geophysical Research Letters*, 43(12):6410–6417.