

TOPICAL REVIEW • **OPEN ACCESS**

Adapting attribution science to the climate extremes of tomorrow

To cite this article: Luke J Harrington and Friederike E L Otto 2018 *Environ. Res. Lett.* **13** 123006

View the [article online](#) for updates and enhancements.

Recent citations

- [Attributable damage liability in a non-linear climate](#)
Luke J. Harrington and Friederike E. L. Otto
- [Embracing the complexity of extreme weather events when quantifying their likelihood of recurrence in a warming world](#)
Luke J Harrington *et al*



TOPICAL REVIEW

OPEN ACCESS

RECEIVED
7 September 2018

REVISED
28 November 2018

ACCEPTED FOR PUBLICATION
29 November 2018

PUBLISHED
18 December 2018

Original content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



Adapting attribution science to the climate extremes of tomorrow

Luke J Harrington¹ and Friederike E L Otto

Environmental Change Institute, University of Oxford, South Parks Road, Oxford, OX1 3QY, United Kingdom

¹ Author to whom any correspondence should be addressed.E-mail: luke.harrington@ouce.ox.ac.uk**Keywords:** climate change, attribution, extreme events, future projections, saturation, adaptationSupplementary material for this article is available [online](#)**Abstract**

Increasing risks of extreme weather events are the most noticeable and damaging manifestation of anthropogenic climate change. In the aftermath of an extreme event, policymakers are often called upon to make timely and sensitive decisions about rebuilding and managing present and future risks. Information regarding whether, where, and how present day and future risks are changing is needed to adequately inform these decisions. But this information is often not available on the temporal and spatial scales decisions are made. In particular, decision makers require information about both historical changes and plausible future changes in the severity and frequency of extreme weather in a seamless way. However, applying the same methods from event attribution to future projections by defining events based on present day frequency of occurrence leads to potentially misleading estimates of future changes in a warmer climate. We demonstrate that this is fundamentally a consequence of risk ratios saturating at different values. This study investigates the circumstances under which present-day attribution frameworks become ill-suited for characterising changes in future extremes, before discussing what alternative frameworks may be more useful to inform stakeholders about what additional risks from extreme weather events will emerge in a warmer world.

1. Introduction

Extreme weather events and their impacts are the main societal manifestation of a changing climate, and often the primary mechanism through which climate-related losses and damages occur. With increasing computational capabilities and a spur of methodological development in recent years, a lot of research has focussed on projecting, understanding and attributing extreme weather events (Seneviratne and Zwiers 2015). The societal need to understand current and future risks from extreme weather events is an important motivation behind the research, hence identifying meteorological events that cause major impacts is key to providing the most useful information.

Defining an event in such an impact-relevant way however is far from straightforward. In the science of extreme event attribution in particular, the use of changing risks of an event occurring due to changing external forcings has been established by calculating risk ratios between different forcing scenarios where the risk of the event occurring under present day

conditions is often taken as the event definition. For example, Philip *et al* (2017) identified the lack of rainfall over a large area of Ethiopia during February–September 2015 to be a 1-in-260 year event, as estimated from fitting a Generalized Extreme Value distribution to historical observations, and then used this 1-in-260 year low-rainfall anomaly in subsequent model analysis as definition of the event. Other analyses identify event thresholds for single variables ranging from the frequently used 1-in-100 year event (Schaller *et al* 2016) to assessing events so exceptionally rare that the lower bound of their return time was an estimated 1-in-5000 years (Risser and Wehner 2017, van Oldenborgh *et al* 2017). The practice of using return periods instead of observed magnitudes of events allows the use of climate models that are able to reproduce observed distributions but are biased with respect to representing the intensity of rainfall or other variables. Such practice is particularly important when event definitions include more than one variable, since bias correction of the model often leads to the compromise

of internal consistency between physical variables (Sippel *et al* 2016).

Further to these event-specific definitions framed by the actual impacts from an observed event, several standardised indices to estimate changes in extreme weather in a changing climate are also based on the probability of occurrence in a fixed baseline period (Zhang *et al* 2011). Examples are TX90p: representing the number of days each year with maximum temperatures above the location- and month-specific 90th percentile; or R95p, representing annual rainfall from all days which exceed the 95th percentile—all such metrics are commonly defined using baseline periods of 1961–1990. Other studies simply analyse how the likelihood of exceeding an event corresponding to the N th percentile in a pre-industrial climate might change under future warming scenarios (Fischer and Knutti 2015).

In the following analysis, we will highlight some concerns which arise when simply transposing event-specific attribution frameworks to infer probabilistic changes to climate extremes in the future, as well as outlining ways to avoid misinterpretation.

2. The saturation of risk ratios

In the increasingly common method of calculating risk ratios—defined as the ratio between the likelihood of an event occurring in a baseline climate (e.g. pre-industrial, or present-day climate), p_0 , and the likelihood of the same event happening in a different climate (present-day, or future climate, respectively), p_1 —the uncertainty in any quantification of the ratio p_1/p_0 depends strongly on the return time of the event chosen, as well as the type of distribution the event variable follows. This is primarily the result of greater uncertainty in quantifying the very ends of the distribution tails for historical weather records which often span less than a century.

There is however another aspect arising solely from the definition of the risk ratio. With the likelihood or probability of occurrence, p , being equal to the inverse of the return period, n , this means that there is a maximum possible risk ratio depending on these return times. In fact, the maximum risk ratio (RR_{MAX}) is equal to the return period of an event in the baseline distribution.

In figure 1 we use Gaussian distributions that could represent large scale temperature changes (Hannart *et al* 2016) to illustrate the implications of this simple fact. Considering an event (vertical lines in figure 1(a)) with an n -year return period with respect to a counterfactual distribution (blue) the likelihood of such an event to occur (p_0) is, by definition:

$$p_0 = \frac{1}{n_0}$$

with the probability of occurrence in the factual distribution (red), p_1 , being a function of the distributional shift as a whole. The corresponding risk ratio, RR , is defined as:

$$RR = \frac{p_1}{p_0} = \frac{n_0}{n_1}.$$


We now explore how the risk ratio for three events, defined with a 1-in-10, 1-in-100 and 1-in-1000 year return period, evolves under a warming climate. To consider in particular how the boundedness or saturation of the risk ratio depends on the overall warming and to explore an emergence into an unknown climate, we slowly increase the mean of the distribution to a six degree difference, where local temperatures increase at a rate equivalent to global-mean warming. Eventually, one finds the factual distribution has moved completely to the right of the event thresholds of interest, at which point the maximum possible ‘factual’ probability reaches a saturation value of 1,

$$\begin{aligned} p_{1,MAX} &\equiv 1 = n_1 \\ \therefore RR_{MAX} &= \frac{p_{1,MAX}}{p_0} \\ &= \frac{1}{1/n_0} \\ \therefore RR_{MAX} &= n_0. \end{aligned}$$

Therefore, the maximum possible risk ratio that can be obtained for a given event threshold is equal to the return period of that event in the counterfactual world. For example, the maximum possible RR that can be witnessed for 1-in-1000 day extreme temperatures (or 99.9th percentile event) is 1000.

2.1. Properties of risk ratio evolution in a Gaussian distribution

Figure 1(a) shows the emergence from the present-day (blue) into an unprecedented climate (red) for a hypothetical Gaussian distribution. Figures 1(b) and (c) show how the risk ratio changes with warming for three event thresholds with return periods in the present-day of 1-in-10 (black), 1-in-100 (brown) and 1-in-1000 (cyan) years. Figure 1(b) presents global warming of up to two degrees while figure 1(c) presents an extended warming period of up to six degrees. Focusing only on the modest warming (figure 1(b)), risk ratios appear to be climbing exponentially for all event thresholds with the most rare events experiencing the fastest increases in frequency, consistent with previous research (Fischer and Knutti 2015, Kharin *et al* 2018). But with further warming, figure 1(c) reveals that a saturation of risk ratios occurs, with each event threshold converging to different risk ratio maxima which are proportional to the return period of the event in the baseline climate. Since the emergence of heat-related signal-to-noise ratios following a given increment of global warming can differ substantially between tropical and extratropical regions (Harrington *et al* 2017), a spatially



explicit representation of future changes in risk based on a present-day baseline will capture different parts of these sigmoids presented in figure 1(c), and thus lead to a distorted interpretation of which regions experience the most severe changes in extreme temperatures.

2.2. Properties of risk ratio evolution in a Gumbel distribution

The results from the Gaussian example in section 2.1 are particularly informative, as they characterise many different types of extreme heat events which occur over large spatio-temporal timescales, such as short heatwaves which span continents (Fischer *et al* 2013, Fischer and Knutti 2015) or extremely hot summers or years for a specific location (Christidis *et al* 2015, King *et al* 2015). However many other types of extreme events can be more closely represented with extreme value statistics (Sippel *et al* 2015, Kharin *et al* 2018). These heavier-tailed distributions have represented the statistical properties of extreme rainfall events particularly well in previous extreme event analyses (Wolski *et al* 2014, Rosier *et al* 2015, Wiel *et al* 2017,

Otto *et al* 2018, Philip *et al* 2018), and thus also warrant examination.

Figure 2(a) shows the warming-driven emergence into an unprecedented climate for an idealised Gumbel distribution, with figures 2(b) and (c) showing how the risk ratio changes with warming for the same three event thresholds as figure 1, defined as return periods in the present-day of 1-in-10 (black), 1-in-100 (brown) and 1-in-1000 (cyan) years.

When compared against the corresponding Gaussian risk ratio profiles, several differences emerge for the Gumbel case study. The first and most obvious is that initial changes in risk ratio are much slower for the heavy-tailed example than for the Gaussian case, largely due to the 1-in- N year event thresholds being positioned proportionally further in the tail by definition.

A second point of difference: in the Gaussian example, figures 1(b) and (c) show the risk ratios for the most extreme event thresholds to always remain higher than corresponding event thresholds which are less rare, irrespective of the increment of warming which had occurred. For example, after the first several

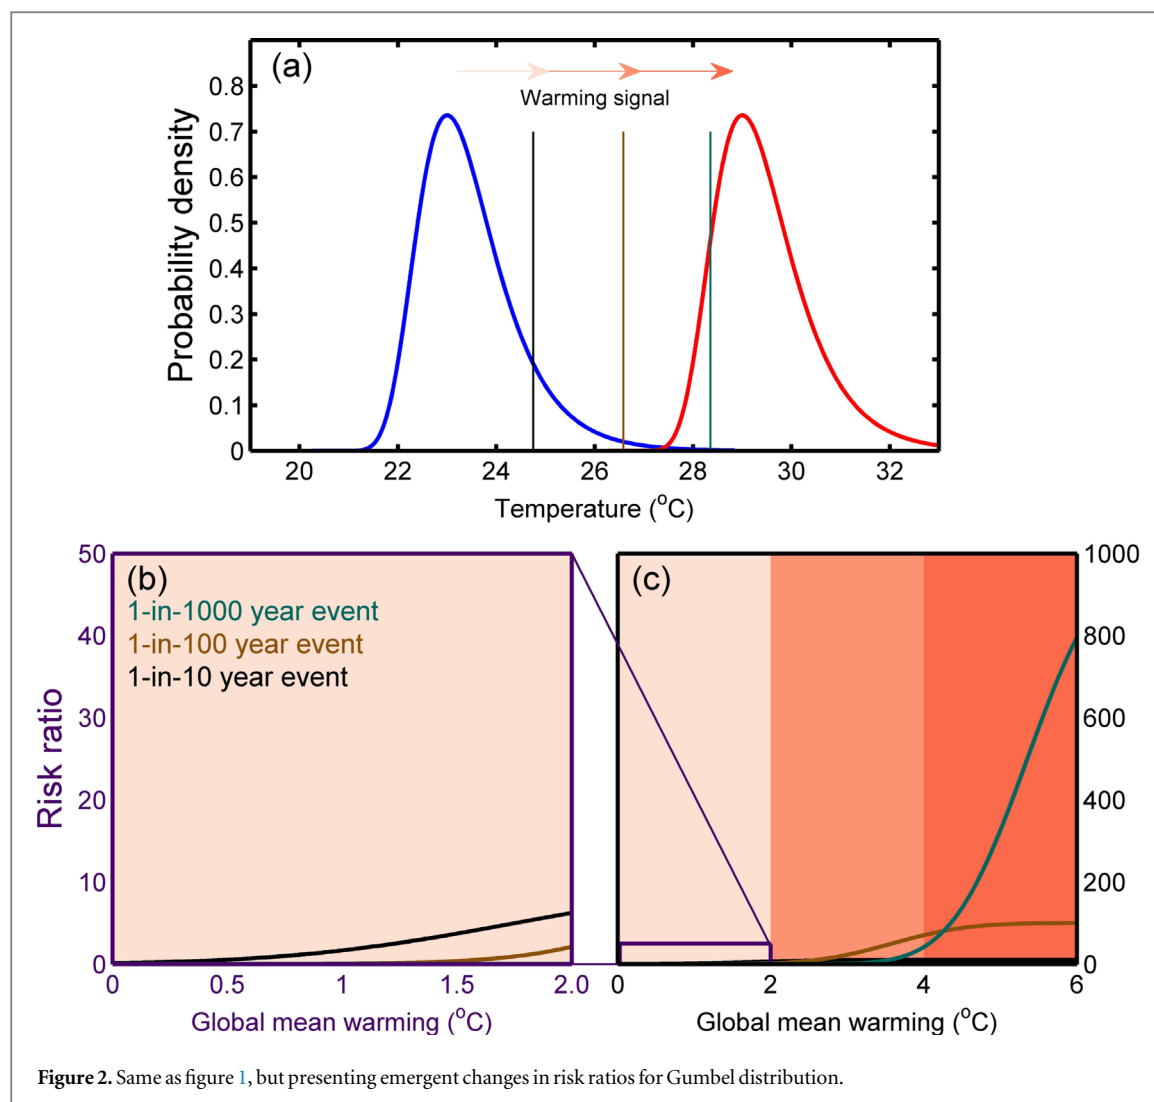


Figure 2. Same as figure 1, but presenting emergent changes in risk ratios for Gumbel distribution.

degrees of warming (figure 1(b)), much more rapid RR increases occur for the 1-in-1000 year event than for the 1-in-10 year event threshold (due to the different thresholds converging to RR maxima (figure 1(c)) which differ by two orders of magnitude). However, this general principle of the most extreme events exhibiting the highest RRs with further warming does not hold for heavy-tailed distributions. Instead, we see in figure 2(b) the smallest event threshold (black line) exhibiting the largest risk ratios for the earlier warming increments, before being overtaken by the 1-in-100 year event RR once warming exceeds two degrees. Similarly in figure 2(c), the risk ratio for the 1-in-100 year event threshold outpaces that of the 1-in-1000 year event until after four degrees of warming, at which point the risk ratio rapidly approaches its saturation value. Such differences reflect the often subtle consequences of varying second- and third-order moments between different statistical distributions (Sardeshmukh *et al* 2015).

The fact that risk ratios eventually saturate at a maximum value equal to the ‘counterfactual’ return period of the event threshold chosen remains true irrespective of the shape of the PDF. However, the relative

pathways by which this saturation occurs are clearly much less intuitive and thus require further investigation.

2.3. Relevance in the context of existing event attribution literature

The saturation effect described above becomes relevant in the case when the event is defined with respect to the return time of the event in the past climate. For the majority of attribution studies published to date, the issue of risk ratio saturation is not a significant problem, as the event is usually defined with respect to a present day return period (to reflect the observed impacts of the event). However, for those studies which consider changing extreme event frequencies with an event defined with respect to the pre-industrial climate (Lewis and Karoly 2013), these threshold-specific upper bounds on plausible risk ratios can distort the interpretation of results.

In addition scientists and stakeholders alike focus increasingly on assessing changing risks of extreme events across timescales (Otto *et al* 2018): not only comparing a counterfactual world with the present day climate but extending studies to assess how

hazards change in various future warming scenarios (Fischer and Knutti 2015). In these cases, when defining the event based on present day return times and calculating the risk ratio under different future scenarios, the saturation effect can influence the calculated change in risk for the future.

Furthermore, with rare events and in data poor regions, the uncertainty bounds on the observed event threshold can be substantive, with ranges from a 1-in-50 to a 1-in-500 year event for the earlier example of Philip *et al* (2017). In this case the RR_{MAX} that could be found will also span orders of magnitude, rendering the interpretation of risk ratio confidence intervals from an attribution analysis rather difficult. For example, Pall *et al* (2017) demonstrates that the lower uncertainty bound of a risk ratio remains similar across a wide range of counterfactual event thresholds, despite the best-guess RR appearing to accelerate in line with expected changes to the corresponding RR_{MAX} —the reason for the lower uncertainty bound remaining stable is therefore likely a result of widening uncertainty bounds compensating for the threshold-driven increases in RR.

When saturation of risk ratios at the tail of the distribution occurs, this does not mean that the impacts of an additional warming become negligible, rather the metric itself begins to no longer represent a good proxy for changes to extremes. In light of different rates of emergent temperature change for different parts of the world, this saturation effect might also lead to people arguing, incorrectly, that high latitudes with slower rates of emergence and thus steeper *initial* changes in risk ratios (further away from saturation) would be the greater beneficiaries for ambitious mitigation. Distortion effects could also manifest in studies designed to address questions like ‘when in the future a doubling of risk will occur’, which could potentially be relevant for questions of liability (Marjanac *et al* 2017, Marjanac and Patton 2018).

While it could be argued that in cases of precipitation events we will never see an emergence of risks that is close to a complete separation of the distributions, examples of such changes do already exist for high-temperature events (Christidis *et al* 2015, Uhe *et al* 2016, Oliver *et al* 2017, Imada *et al* 2018, Knutson *et al* 2018, Walsh *et al* 2018). In the case of European heatwaves, Christidis *et al* (2015) not only analysed an event where the distributions separate strongly but also chose a threshold based on a past record (2003 in this case) thus defining the event in a way for it to be potentially misinterpreted, say if the relative impacts of the first degree of additional global warming (from today) were compared with a second degree of temperature rise. Being an artefact from the definition of the risk ratio, this saturation is most likely not a good representative of changing impacts of anthropogenic climate change in the real world.

3. Understanding the influence of future warming on the ‘extreme events’ of today

While likely a small problem in attribution studies so far, the above shows that assessments of future changes in extremes can be strongly influenced by statistical artefacts in the choice of metric. To be able to identify those it is important to understand how the properties of the RR as well as the corresponding fraction of attributable risk (FAR) - defined as $1/(1-RR)$ - behave as a function of the underlying distribution under future warming.

To obtain such a more systematic understanding of how both risk ratios and FARs evolve with continued warming, we next examine the changes to a range of possible extreme event thresholds for a sequence of specific, warming-induced deviations from a hypothetical statistical distribution which represents a climate variable in the present-day. To more easily link distributional changes with corresponding global warming for the reader, we hereafter associate the shift of a distribution by one standard deviation with one degree of global mean temperature increase from the present day.

3.1. The impact of future warming across plausible event thresholds: Gaussian distribution

We first consider in figure 3(a) the emergence of a Gaussian PDF by six discrete units of half a standard deviation each, hereafter representing 0.5°C increments of global mean warming (thus, the farthest emerged PDF, in red, has experienced a shift of 3 standard deviations following an illustrative global temperature rise of 3°C). Equivalent analyses for smaller warming increments are also presented in the supplementary material which is available online at stacks.iop.org/ERL/13/123006/mmedia and show qualitatively similar results, albeit with less distortion from saturation at risk ratio maxima. Figures 3(b) and (c) show the corresponding RR and FAR for each of these six sequential PDFs compared to the black (present day) PDF, for present-day event threshold spanning return periods of between 1-in-10 and 1-in-10 000 years. In both panels, the grey shaded region denotes the RR/FAR space beyond the relevant maximum value which can be reached for that event threshold.

Figure 3(b) shows a robust exponential increases in risk ratios (which appear linear when plotted with a log scale), both when looking at changes for more extreme event thresholds with respect to a specific PDF (coloured lines), as well as when looking at RR changes for the same event threshold across multiple PDFs with successive increments of warming. This in itself is a significant result, and shows that (assuming statistical properties which resemble a Gaussian distribution, such as for season-average temperatures) any observational uncertainty which exists in

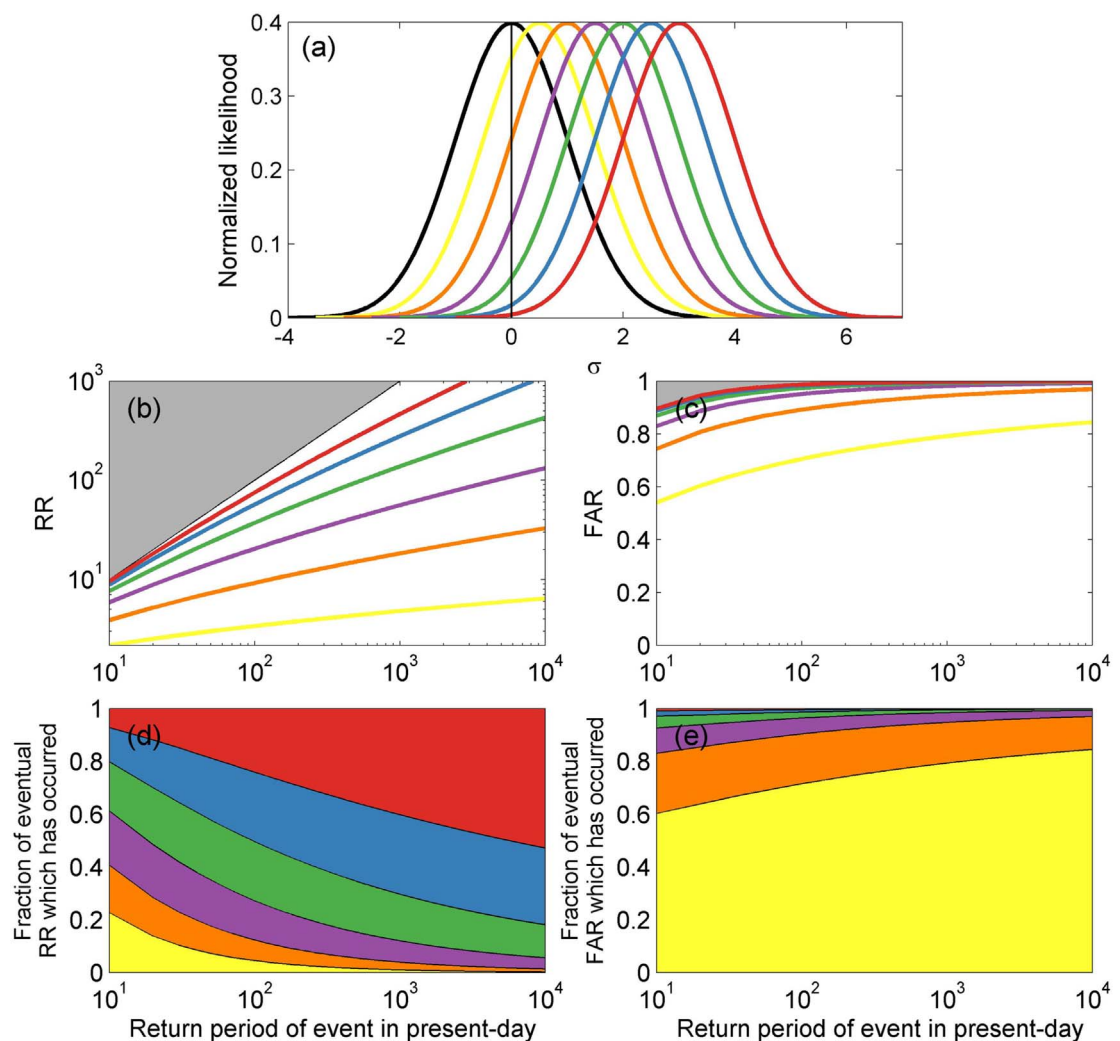


Figure 3. The emergence of absolute and relative RRs and FARs for a warming Gaussian distribution. (a) Schematic of the warming of a 'present-day' Gaussian distribution (black), with each successive colour reflecting a shift in the distribution mean by $+0.5$ standard deviations. (b) Corresponding changes in risk ratio for each of the shifted PDFs, considering all event thresholds which have a return period between 1-in-10 years and 1-in-10 000 years in the present-day climate. Grey shaded areas illustrate risk ratios which exceed the maximum possible value for each event threshold. (c) Same as for panel b, but showing changes in the fraction of attributable risk (or FAR) across all many event thresholds for each of the successively warmer PDFs. Panels (d) and (e) present cumulative area graphs of the fraction of the eventual RR and FAR which occurs after each of the 0.5σ warming increments. In this context, 'eventual' has been defined as the value reached after $+3\sigma$ of warming (the red lines in b and c).

identifying the return period for a damaging extreme event in the present day will translate nonlinearly to uncertainty in how the likelihood of that event will increase into the future.

With respect to the changes to FAR with emergent climate warming, we see in figure 3(c) that only a $+1\sigma$ distribution shift is needed (orange line) before the FAR exceeds 0.8 across all extreme event thresholds, and convergence towards the maximum possible FAR values is observed after much less warming.

While understanding the absolute changes to RR and FAR with continuing warming is a useful exercise, recent analyses have become particularly interested in understanding the *relative* impact of a given increment global temperature increase—that is, how nonlinear the evolution of RR and FAR estimates appear with continued warming. Is it the case that the most rapid and substantive changes to extreme events will happen

within the next 0.5°C of global warming, or will there be a sudden and rapid increase in such changes between a 1.5°C and 2°C world? The answers to these questions are highly relevant for decision makers when thinking about whether to implement moderate or ambitious mitigation policies, and particularly whether the relative benefits of avoided impacts outweigh the costs of implementation for each case.

To confirm whether the evolution of RR and FAR occurs linearly with each increment of warming, as well as whether it emerges similarly across the two metrics, the variations in RR/FAR as a function of event threshold (x-axis) for figures 3(b) and 3(c) are normalised by the changes witnessed for the PDF which has emerged the farthest. Thus we create a new metric (we call 'fraction of eventual RR/FAR') that highlights the relative changes and makes it easy to identify non-linearities. Figures 3(d) and (e) show the

fraction of the eventual risk ratio (or FAR) which emerges during each 0.5°C increment of warming. ‘Eventual’ in this context refers specifically to the values reached after 3°C of warming (red lines in figures 3(b) and (c)).

3.1.1. Relative changes in RR with successive increments of warming

Several important results can be obtained from this metric. Focusing first on figure 3(d), we see that the ‘fraction of eventual RR’ which is occupied by the different colours can differ dramatically when considering different event thresholds along the horizontal axis. For example, when looking at changes to the frequency of a present-day 1-in-10 year event, we find the fractional increases in RR to emerge most within the first 0.5°C increment of warming (yellow) and the smallest changes in RR to occur within the final 0.5°C warming increment (red). But when one looks (along the x-axis) at the relative changes for a present-day 1-in-10 000 year event, we find that less than 10% of the total RR is witnessed after 1.5°C of warming (purple), and during only the final 0.5°C warming increment (red) does the risk ratio increase by the same amount as for the preceding 2.5°C of warming.

When looking at the collective patterns of figure 3(d) over all event thresholds, one finds that each sequential increment of warming (colour band) leads to a larger change in risk ratio than the preceding increment of warming. However, the extent of this nonlinear rise in RR with each further warming increment is systematically suppressed as one shifts from considering the most severe event thresholds to more moderate event thresholds: the reason for this suppression is fundamentally due to the fact that those less severe event thresholds reach saturation at RR_{MAX} earlier (as discussed in section 2).

3.1.2. Relative changes in FAR with successive increments of warming

For the case of FAR emergence, figure 3(e) reveals the pattern of relative increases in the FAR differ significantly from the corresponding risk ratio patterns in figure 3(d). Indeed, more than 80% of the eventual FAR is reached after only 1°C of warming (yellow and orange bands), irrespective of the event threshold considered. Moreover, a larger fraction of these changes in FAR happens within the first 0.5°C increment of warming for the most extreme event thresholds.

These patterns are an understandable artefact of the nonlinear statistical relationship between RRs and FARs, whereby significant risk ratios (>10) which are found to emerge for the more ‘extreme’ event thresholds after a large distributional shift, simply translate to variations in FAR between 0.9 and 1. The formulation of FAR means the saturation of values close to 1 occurs much more rapidly when compared with the warming required for the risk ratios to saturate near

their corresponding maxima. While a consequence of simple statistics, this difference between the warming required for FARs to saturate and risk ratios to saturate could lead to potentially misleading interpretations, especially if a decision maker were to consider the FAR metric in isolation when deducing the relative detriment of a given increment of future global warming.

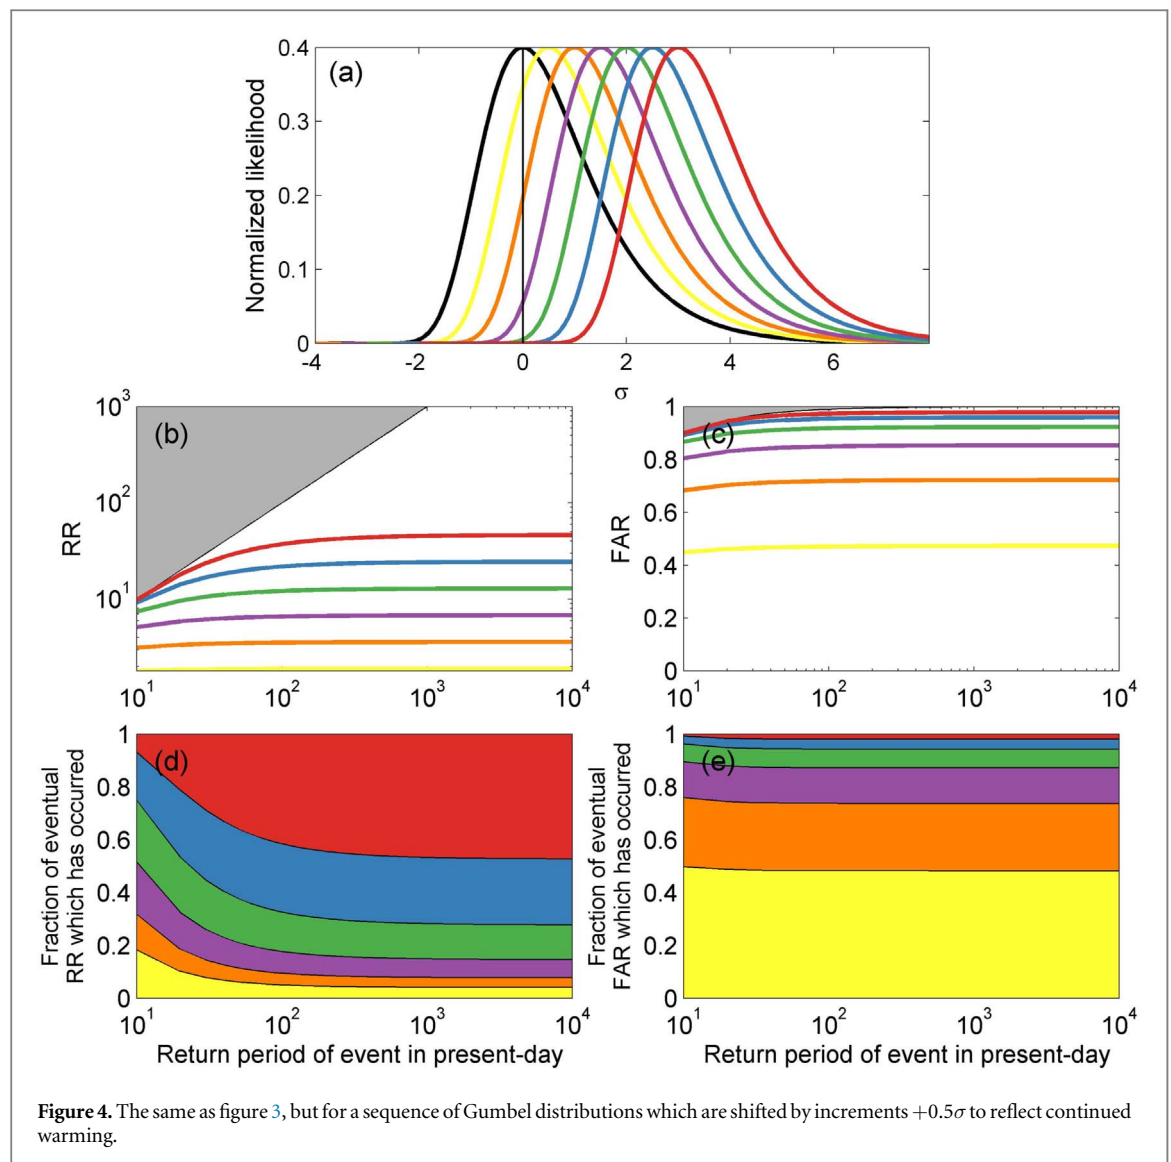
3.2. The impact of future warming across plausible event thresholds: Gumbel distribution

When considering the warming-driven evolution of RRs and FARs for a Gumbel distribution in figure 4, several points of difference emerge relative to the Gaussian-based results seen in figure 3. As already shown in figures 1 and 2, 4(c) reveals that shifting a Gumbel distribution by a certain number of standard deviations as a result of global warming results in much lower risk ratios when compared against an equivalent shift in the Gaussian case. Such a result is largely unsurprising and reflects the substantially larger distribution tail which exists in a Gumbel distribution.

However, a more interesting consequence of the Gumbel distribution is that both RRs and FARs remain constant for all event thresholds which have a present day return period exceeding 1-in-100 years following a given increment of warming. Such a result is in stark contrast to the Gaussian example, highlighting a much smaller sensitivity to the exact event definition for extreme events.

When looking at the relative influence of each successive warming-induced shift of the Gumbel distribution, results generally reveal accelerating increases in each subsequent risk ratio (red shading larger than yellow in figure 4(d)) and saturating changes to each subsequent FAR (yellow shading larger than red in figure 4(e)). Such conclusions are consistent with the Gaussian case for those highest (1-in-10 000 year) event thresholds unaffected by RR saturation: each successive increment of future warming beyond the present-day leads to successively larger increases in the frequency of extremes, but smaller increases in the corresponding FAR.

A key result from this comparison is that changes in extremes that are occurring more frequent than 1-in-100 years in the present day, independent of the distribution, are potentially subject to misinterpretation in the future, owing to risk ratios approaching saturation after subsequent warming. Moreover, supplementary figures 3 and 4 show these conclusions to be robust, even if increases in temperature variability (and hence a widening of the PDF) were to accompany the warming-driven shifts in the distribution mean (Holmes *et al* 2016, Bathiany *et al* 2018, Guirguis *et al* 2018). This is because increasing variability alongside a warming-driven shift in the distribution overall positions a larger fraction of the PDF in exceedance of present-day event thresholds.



4. Discussion

4.1. Is there any suitable metric of aggregate changes to future climate extremes?

Section 3 has highlighted the numerous ways in which the changes to RRs and FARs with continuing warming can exhibit substantive variability and nonlinearity, depending on both the extreme event threshold which is chosen as well as the range of future warming which is considered and the shape of the distribution.


These sensitivities demonstrate the fundamental limitations which exist when trying to quantify aggregate changes to extreme weather in an exceedingly warmer world. It is impossible to perfectly quantify the impacts which result from future events which will exceed the return period of any historical extremes by several orders of magnitude. The best approach has been instead to examine the damages which have emerged from extreme events observed in the real world, assume there to be some coherent (though potentially nonlinear) relationship between event rarity and impacts, and then take these observed impacts

as a proxy to inform assessments of risk from unprecedented extreme events in a future climate which may well be unrecognisable. However, any confidence in speculating about the relationship between impacts associated with observed extreme events, and the impacts of those future events which may be many orders of magnitude more severe, diminishes as this extrapolation in severity increases—such is the case for statistical distributions which separate entirely.

Despite these fundamental constraints, some qualitative guidelines do emerge from the analysis in section 3 to avoid misinterpreting which N -degree increment of warming leads to the most significant changes in extreme events. This includes ensuring a range of multiple event thresholds are analysed and that each of these thresholds would still be ‘rare’ even under the most extreme warming scenario considered.

4.2. Alternatives to the RR or FAR metric when understanding the emergence of future extremes

Examining the relative effect of future warming increments on changing risk ratios for a variety of



different event thresholds is but one of several ways to understand the decision-relevant changes to extreme weather-related events which might occur in a warmer world.

From the perspective of infrastructure planning, a more common approach is to identify intensity thresholds which correspond to specific return periods in the present day and which systems are then designed to cope with. Stormwater drainage systems, for example, might be designed to withstand a 1-in-500 year flooding event based on statistical analysis of historical rainfall records. However, events of equal intensity might become only a 1-in-100 year event after global temperatures rise by several degrees. Understanding planning horizons and what changes to infrastructure are needed therefore depends on the potential changes in these return periods, and the associated impacts that a several-fold increase in the frequency of system failures might bring.

Another common example of a specific extreme event threshold carrying added significance from a decision-making perspective relates to property insurance for locations which may be periodically exposed to flooding inundation events (either of riverine or coastal origin). Local planners commonly obtain flood exposure maps for residential properties located near rivers or coastlines, showing the annual probability of inundation occurring on their property. If flooding is determined to occur with a return period of, say, 1-in-20 years, that land may then be deemed uninsurable (Storey and Noy 2018).

In this context, a better way to understand the relevant changes to future extremes is to consider the

additional warming required for present-day return periods to *become* as common as the return period of interest (like a 1-in-20 year event). That way, the map of present-day exposure probabilities can be simply converted into a map showing the added warming required before that land is no longer insurable. Figure 5 presents such an example for both a Gaussian and GEV distribution, illustrating how much of a warming-driven shift in each distribution is needed for many different present-day event thresholds to converge to decision-relevant return periods in the future. Consistent with the results found in figures 1–4, the relative warming required for changes to occur in a Gaussian distribution is much less than for a GEV (Gumbel) distribution: for example, an estimated 1.5 standard deviations of local change is needed for a 1-in-100 year event to become a 1-in-5 year event under Gaussian assumptions, but about 2.5 standard deviations is required for an equivalent change to occur in the Gumbel example.

4.3. Computational barriers to understanding changes in distribution tails with continued warming

When attempting to understand the anthropogenic contribution to extreme weather changes in the present day, or the warming-driven changes to extreme events in the future, computational expense has necessitated a trade-off for those climate model experiments which have been developed to date. One option has been to simulate many, many model runs for only a small number of discrete model worlds (such as the climate in the present day, the climate of

today in the absence of human influence, or the climate of a 1.5 °C or 2 °C world): examples include the *Weather@Home* (Guillod *et al* 2017), HAPPI (Mitchell *et al* 2017) and C20C + projects (Angéllil *et al* 2017). The other option is to generate long transient simulations of evolving future simulations which cross multiple temperature thresholds, and then pool model years together to generate the sample sizes needed for analysing distribution tails (James *et al* 2017): examples include the CMIP5 (King and Karoly 2017) and CESM-LENS (Kay *et al* 2015) model frameworks.

Unfortunately, accurately quantifying the future changes to climate extremes, both across a wide spectrum of plausible event thresholds at the distribution tail and for a range of continuous warming increments, requires not only transient simulations passing many temperature thresholds but also the repetition of these long-run experiments many thousands of times. While this requires substantive increases in computational resources, the results demonstrated in section 3 emphasise the importance of such work to understand the characteristics of changing extreme weather risks across multiple future temperature horizons.

4.4. Recommendations for understanding specific extreme events from the present-day in the context of a warming world

In light of the results found under the idealised examples in sections 2 and 3, several key sensitivity tests have emerged for researchers to consider, when analysing future changes to the likelihood of witnessing an ‘extreme event’ as defined from the present-day climate:

- (1) Are the risk ratios remaining constant over a range of event thresholds? If not, then it is likely that arbitrarily changing the event threshold to a rarer return period will also arbitrarily increase the risk ratio (figure 3(b)). Similar questions apply when considering possible differences in the risk ratio for an observed extreme event under future warming—particularly once the large uncertainty that often exists in quantifying the ‘observed’ return period of the event is accounted for. One might still want to proceed with the analysis, but the dependency of the RR from the threshold will be as important a result as the RR itself.
- (2) Are the changes in risk ratio over multiple warming increments consistent for multiple event thresholds? If not, then the researcher needs to explicitly state that simply choosing a different event threshold will lead to a different interpretation of which N -degree increment of warming generates the most ‘severe’ changes in extreme event frequency. This is particularly pertinent in the context of comparing temperature increments

which carry specific policy-relevance, such as the differences between a world under 1.5 °C or 2 °C of global warming.

- (3) Are the *relative increases* in risk ratio with each successive increment of warming larger than for the previous warming increment? If not, then it is likely that the probability of exceeding the specified event threshold has become too high, such that interpretations of the relative change in risk ratio with warming are becoming skewed by partial risk ratio saturation (see figures 1(c), 2(c), 3(d) and 4(d)). When the probability of exceeding any ‘event threshold’ becomes sufficiently common—the authors propose larger than 0.2—then the value of metrics like the risk ratio, particularly as a tool to inform decision makers, begins to break down.

5. Conclusions

The framework of present-day event attribution studies currently involves defining an event threshold as close as possible to some observed (and often damaging) incident in the real world, so to most closely characterise the impacts that happened as a result. If one then utilises that same event definition to discuss future changes in climate extremes, such an approach provides a singular estimate of the changes in frequency for one specific threshold, with those results then being used as a proxy to infer the collective changes in frequency of all plausible extreme events which may be seen in the future.

This analysis has highlighted multiple concerns which arise when simply transposing event-specific attribution frameworks to infer probabilistic changes to climate extremes in the future. If the distribution of the climate variable of interest exhibits Gaussian statistical properties, then the risk ratios under future warming will increase purely as a consequence of selecting a higher event threshold to examine (figure 1). Second, selecting a single event threshold to infer changes in the frequency of occurrence in the future will lead to a very specific interpretation of *when* the fastest changes in risk ratio will occur with each successive warming increment (figure 3). Third, each of these issues which apply for RR interpretations are amplified further when one considers corresponding changes to the FAR.

While the case for heavier-tailed PDFs like the Gumbel distribution is more complex, particularly in terms of the sigmoidal evolution as risk ratios converge towards their saturation value (figure 2), the capacity to accurately interpret the relative impact of a specific N -degree increment of warming on changing event frequency remains equally problematic.

There are inherent differences in the questions being asked when attributing probabilistic changes to extreme events in the present day as a result of human

influences, versus understanding aggregate changes in the frequency of all potential extreme weather events in an increasingly warmer future. The frameworks for analysing such future projections need to reflect these differences. We have suggested several alternative frameworks which provide clearer practical relevance for decision-makers (figure 5), though these carry the burden of proportionally higher computational demands.

Once the future distribution of weather starts to deviate significantly from anything recognisable in the historical record, the relevance of links between events previously considered ‘extreme’, and their corresponding impacts, increasingly diminishes. Thus, not only will unprecedented changes in local climate carry significant impacts, but the capability of scientists to inform and prepare decision makers for these changes will decrease as well. This analysis consequently provides yet another incentive for meaningful and sustained reductions in carbon emissions over the coming decades.

Acknowledgments

The authors would like to thank Dave Frame and Suzanne Rosier for multiple useful discussions related to the manuscript. LJH acknowledges support from the MaRIUS project: Managing the Risks, Impacts and Uncertainties of droughts and water Scarcity, funded by the Natural Environment Research Council (NE/L010364/1).

ORCID iDs

Luke J Harrington  <https://orcid.org/0000-0002-1699-6119>

Friederike E L Otto  <https://orcid.org/0000-0001-8166-5917>

References

- Angélil O, Stone D, Wehner M, Paciorek C J, Krishnan H and Collins W 2017 An independent assessment of anthropogenic attribution statements for recent extreme temperature and rainfall events *J. Clim.* **30** 5–16
- Bathiany S, Dakos V, Scheffer M and Lenton T M 2018 Climate models predict increasing temperature variability in poor countries *Sci. Adv.* **4** eaar5809
- Christidis N, Jones G S and Stott P A 2015 Dramatically increasing chance of extremely hot summers since the 2003 European heatwave *Nat. Clim. Change* **5** 46–50
- Fischer E M, Beyerle U and Knutti R 2013 Robust spatially aggregated projections of climate extremes *Nat. Clim. Change* **3** 1033–8
- Fischer E M and Knutti R 2015 Anthropogenic contribution to global occurrence of heavy-precipitation and high-temperature extremes *Nat. Clim. Change* **5** 560–4
- Guillod B P *et al* 2017 weather@home 2: validation of an improved global—regional climate modelling system *Geosci. Model Dev.* **10** 1849–72
- Guirguis K, Gershunov A, Cayan D R and Pierce D W 2018 Heat wave probability in the changing climate of the Southwest US *Clim. Dyn.* **50** 3853–64
- Hannart A, Carrassi A, Bocquet M, Ghil M, Naveau P, Pulido M, Ruiz J and Tandeo P 2016 DADA: data assimilation for the detection and attribution of weather and climate-related events *Clim. Change* **136** 155–74
- Harrington L J, Frame D J, Hawkins E and Joshi M 2017 Seasonal cycles enhance disparities between low- and high-income countries in exposure to monthly temperature emergence with future warming *Environ. Res. Lett.* **12** 114039
- Holmes C R, Woollings T, Hawkins E and de Vries H 2016 Robust future changes in temperature variability under greenhouse gas forcing and the relationship with thermal advection *J. Clim.* **29** 2221–36
- Imada Y, Shiogama H, Takahashi C, Watanabe M, Mori M, Kamae Y and Maeda S 2018 Climate Change Increased the Likelihood of the 2016 Heat Extremes in Asia *Bull. Am. Meteor. Soc.* **99** S97–101
- James R, Washington R, Schleussner C-F, Rogelj J and Conway D 2017 Characterizing half-a-degree difference: a review of methods for identifying regional climate responses to global warming targets *WTREs Clim. Change* **8** e457
- Kay J E *et al* 2015 The community earth system model (cesm) large ensemble project: A community resource for studying climate change in the presence of internal climate variability *Bull. Amer. Meteor. Soc.* **96** 1333–49
- Kharin V V, Flato G M, Zhang X, Gillett N P, Zwiers F and Anderson K J 2018 Risks from climate extremes change differently from 1.5 °C to 2.0 °C depending on rarity *Earth's Future* **6** 704–15
- King A D and Karoly D 2017 Climate extremes in Europe at 1.5 and 2 degrees of global warming *Environ. Res. Lett.* **12** 114031
- King A D, Oldenborgh G J, van, Karoly D J, Lewis S C and Cullen H 2015 Attribution of the record high Central England temperature of 2014 to anthropogenic influences *Environ. Res. Lett.* **10** 054002
- Knutson T R, Kam J, Zeng F and Wittenberg A T 2018 CMIP5 model-based assessment of anthropogenic influence on record global warmth during 2016 *Bull. Amer. Meteorol. Soc.* **99** S11–5
- Lewis S C and Karoly D J 2013 Anthropogenic contributions to Australia's record summer temperatures of 2013 *Geophys. Res. Lett.* **40** 3705–9
- Marjanac S and Patton L 2018 Extreme weather event attribution science and climate change litigation: an essential step in the causal chain? *J. Energy Nat. Res. Law* **0** 1–34
- Marjanac S, Patton L and Thornton J 2017 Acts of God, human influence and litigation *Nat. Geosci.* **10** 616–9
- Mitchell D *et al* 2017 Half a degree additional warming, prognosis and projected impacts (HAPPI): background and experimental design *Geosci. Model Dev.* **10** 571–83
- van Oldenborgh G J, Wiel K, van der, Sebastian A, Singh R, Arrighi J, Otto F, Hausteine K, Li S, Vecchi G and Cullen H 2017 Attribution of extreme rainfall from Hurricane Harvey, August 2017 *Environ. Res. Lett.* **12** 124009
- Oliver E C J, Benthuyzen J A, Bindoff N L, Hobday A J, Holbrook N J, Mundy C N and Perkins-Kirkpatrick S E 2017 The unprecedented 2015/16 Tasman sea marine heatwave *Nat. Commun.* **8** 16101
- Otto F E L, Philip S, Kew S, Li S, King A and Cullen H 2018 Attributing high-impact extreme events across timescales—a case study of four different types of events *Clim. Change* **1–14**
- Pall P, Patricola C M, Wehner M F, Stone D A, Paciorek C J and Collins W D 2017 Diagnosing conditional anthropogenic contributions to heavy Colorado rainfall in September 2013 *Weather Clim. Extremes* **17** 1–6
- Philip S *et al* 2018 Attribution analysis of the Ethiopian drought of 2015 *J. Climate* **31** 2465–86
- Philip S, Kew S F, Jan van Oldenborgh G, Aalbers E, Vautard R, Otto F, Hausteine K, Habets F and Singh R 2018 Validation of a rapid attribution of the May/June 2016 flood-inducing

- precipitation in France to climate change *J. Hydrometeorol.* **19** 1881–98
- Risser M D and Wehner M F 2017 Attributable human-induced changes in the likelihood and magnitude of the observed extreme precipitation during hurricane Harvey *Geophys. Res. Lett.* **44** 12457–64
- Rosier S, Dean S, Stuart S, Carey-Smith T, Black M T and Massey N 2015 Extreme rainfall in early July 2014 in Northland, New Zealand—was there an anthropogenic influence? *Bull. Am. Meteorol. Soc.* **96** S136–40
- Sardeshmukh P D, Compo G P and Penland C 2015 Need for caution in interpreting extreme weather statistics *J. Clim.* **28** 9166–87
- Schaller N *et al* 2016 Human influence on climate in the 2014 southern England winter floods and their impacts *Nat. Clim. Change* **6** 627–34 (<http://nature.com/nclimate/journal/vaop/ncurrent/full/nclimate2927.html>)
- Seneviratne S I and Zwiers F W 2015 Attribution and prediction of extreme events: editorial on the special issue *Weather Clim. Extremes* **9** 2–5
- Sippel S, Mitchell D, Black M T, Dittus A J, Harrington L, Schaller N and Otto F E L 2015 Combining large model ensembles with extreme value statistics to improve attribution statements of rare events *Weather Clim. Extremes* **9** 25–35
- Sippel S, Otto F E L, Forkel M, Allen M R, Guillod B P, Heimann M, Reichstein M, Seneviratne S I, Thonicke K and Mahecha M D 2016 A novel bias correction methodology for climate impact simulations *Earth Syst. Dyn.* **7** 71–88
- Storey B and Noy I 2017 Insuring property under climate change *Policy Q.* **13** 68–74
- Uhe P, Otto F E L, Haustein K, van Oldenborgh G J, King A D, Wallom D C H, Allen M R and Cullen H 2016 Comparison of methods: attributing the 2014 record European temperatures to human influences *Geophys. Res. Lett.* **43** 2016GL069568
- Walsh J E *et al* 2018 The high latitude marine heat wave of 2016 and its impacts on Alaska *Bull. Am. Meteorol. Soc.* **99** S39–43
- Wiel K, van der Kapnick S B, Oldenborgh G J, van Whan K, Philip S, Vecchi G A, Singh R K, Arrighi J and Cullen H 2017 Rapid attribution of the August 2016 flood-inducing extreme precipitation in south Louisiana to climate change *Hydrol. Earth Syst. Sci.* **21** 897–921
- Wolski P, Stone D, Tadross M, Wehner M and Hewitson B 2014 Attribution of floods in the Okavango basin, Southern Africa *J. Hydrol.* **511** 350–8
- Zhang X, Alexander L, Hegerl G C, Jones P, Tank A K, Peterson T C, Trewin B and Zwiers F W 2011 Indices for monitoring changes in extremes based on daily temperature and precipitation data *Wiley Interdiscip. Rev.: Clim. Change* **2** 851–70